

US010126977B2

(12) **United States Patent**
Sinha

(10) **Patent No.:** **US 10,126,977 B2**
(45) **Date of Patent:** ***Nov. 13, 2018**

(54) **SYNCHRONIZING SELECTED PORTIONS OF DATA IN A STORAGE MANAGEMENT SYSTEM**

(56) **References Cited**

(71) Applicant: **Commvault Systems, Inc.**, Tinton Falls, NJ (US)

4,084,231 A 4/1978 Capozzi et al.
4,267,568 A 5/1981 Dechant et al.

(Continued)

(72) Inventor: **Prosenjit Sinha**, Manalapan, NJ (US)

FOREIGN PATENT DOCUMENTS

(73) Assignee: **Commvault Systems, Inc.**, Tinton Falls, NJ (US)

EP 0259912 3/1988
EP 0405926 1/1991

(Continued)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

OTHER PUBLICATIONS

This patent is subject to a terminal disclaimer.

Arneson, "Mass Storage Archiving in Network Environments" IEEE, 1998, pp. 45-50.

(Continued)

(21) Appl. No.: **15/893,281**

Primary Examiner — Baboucarr Faal

(22) Filed: **Feb. 9, 2018**

(74) *Attorney, Agent, or Firm* — Knobbe, Martens, Olson & Bear, LLP

(65) **Prior Publication Data**

(57) **ABSTRACT**

US 2018/0232149 A1 Aug. 16, 2018

Disclosed methods and systems leverage resources in a storage management system to partially synchronize primary data files based on synchronizing selected portions thereof without regard to changes that may be occurring in other non-synchronized portions. Accordingly, a number of primary data files may be partially synchronized by synchronizing designated portions thereof via auto-restore operations from backup data. This approach relies on storage management resources to designate portions of source data that is to be kept synchronized across any number of targets; detect changes to the designated portions; back up changes to secondary storage; and distribute the changes from secondary storage to the associated targets, with minimal impact to the primary data environment. The approach may be mutually applied, so that changes in any one of an associated group of source data files may be likewise detected, backed up, and distributed to the other members of the group.

Related U.S. Application Data

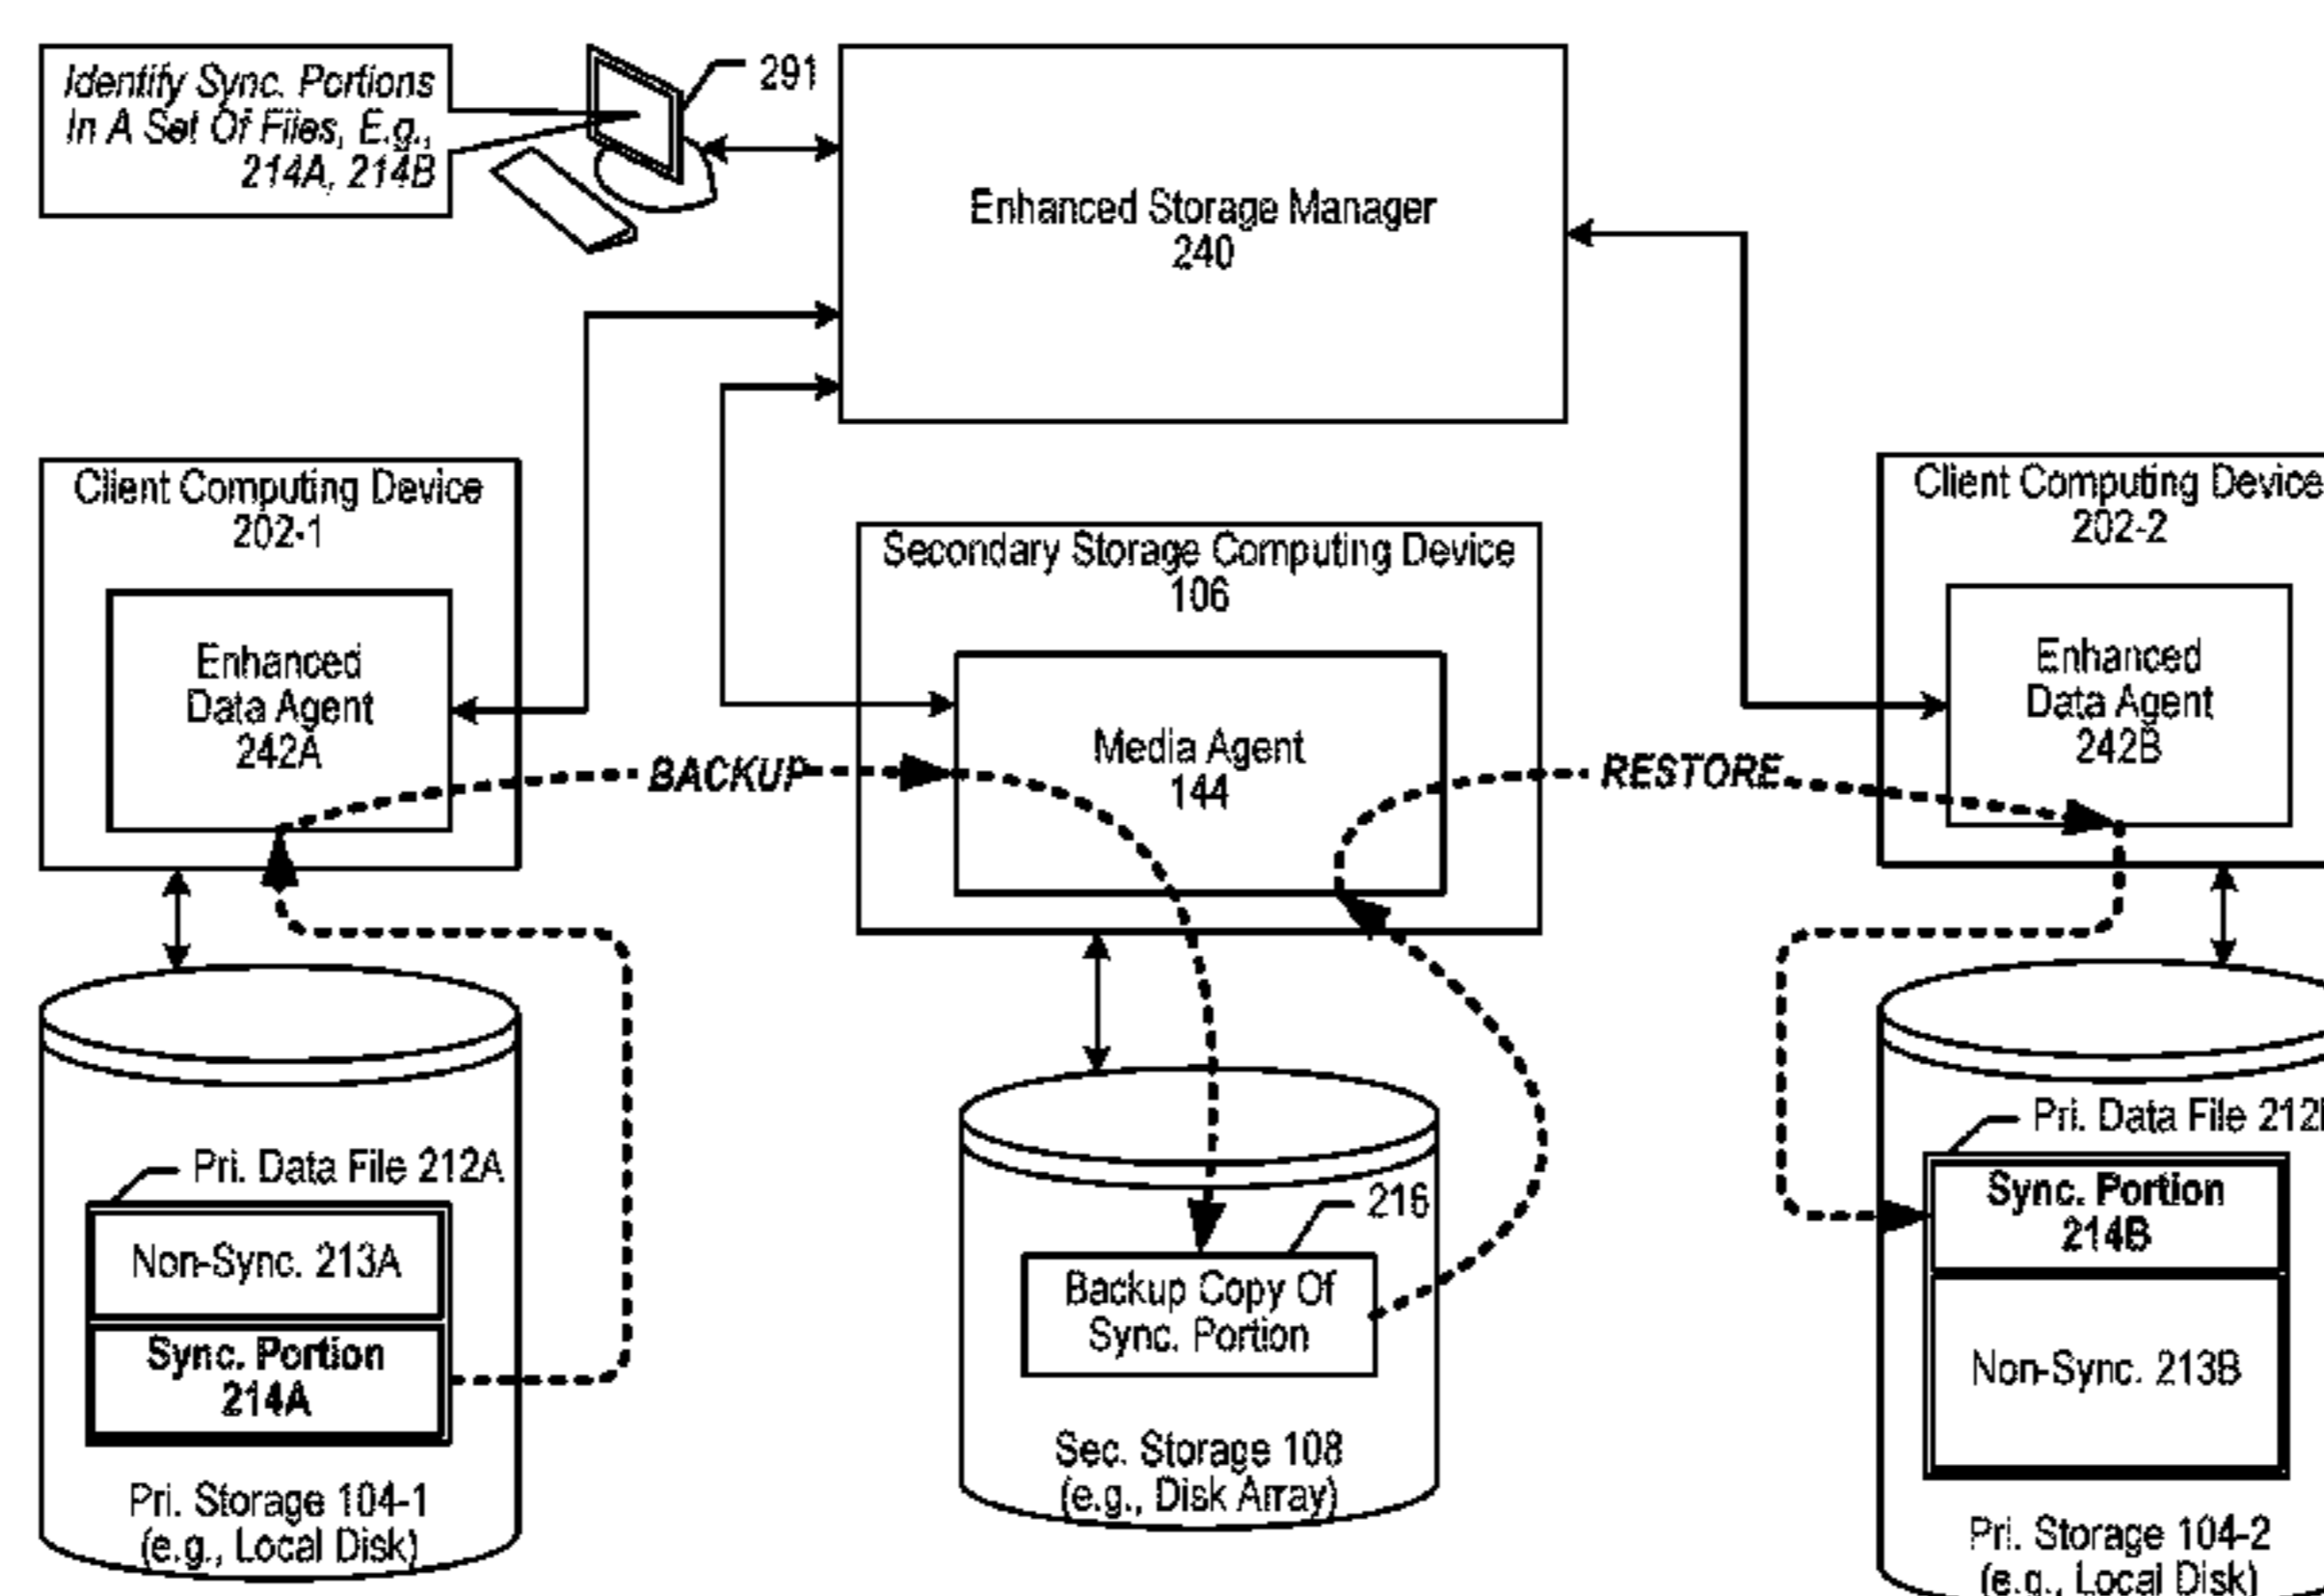
(63) Continuation of application No. 15/412,645, filed on Jan. 23, 2017, now Pat. No. 9,928,005, which is a (Continued)

(51) **Int. Cl.**
G06F 12/00 (2006.01)
G06F 3/06 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 3/065** (2013.01); **G06F 3/067** (2013.01); **G06F 3/0611** (2013.01); **G06F 3/0653** (2013.01)

(58) **Field of Classification Search**
None
See application file for complete search history.

20 Claims, 12 Drawing Sheets



Storage Mgmt. System 200 For Partially Synchronizing Primary Data Files Based On Synchronizing Portions Thereof

Related U.S. Application Data

continuation of application No. 14/601,053, filed on Jan. 20, 2015, now Pat. No. 9,588,849.

(56)

References Cited

U.S. PATENT DOCUMENTS

4,283,787 A 8/1981 Chambers
 4,417,321 A 11/1983 Chang et al.
 4,641,274 A 2/1987 Swank
 4,654,819 A 3/1987 Stiffler et al.
 4,686,620 A 8/1987 Ng
 4,912,637 A 3/1990 Sheedy et al.
 4,995,035 A 2/1991 Cole et al.
 5,005,122 A 4/1991 Griffin et al.
 5,093,912 A 3/1992 Dong et al.
 5,133,065 A 7/1992 Cheffetz et al.
 5,193,154 A 3/1993 Kitajima et al.
 5,212,772 A 5/1993 Masters
 5,226,157 A 7/1993 Nakano et al.
 5,239,647 A 8/1993 Anglin et al.
 5,241,668 A 8/1993 Eastridge et al.
 5,241,670 A 8/1993 Eastridge et al.
 5,276,860 A 1/1994 Fortier et al.
 5,276,867 A 1/1994 Kenley et al.
 5,287,500 A 2/1994 Stoppani, Jr.
 5,301,286 A 4/1994 Rajani
 5,321,816 A 6/1994 Rogan et al.
 5,347,653 A 9/1994 Flynn et al.
 5,410,700 A 4/1995 Fecteau et al.
 5,420,996 A 5/1995 Aoyagi
 5,448,718 A 9/1995 Cohn et al.
 5,454,099 A 9/1995 Myers et al.
 5,559,991 A 9/1996 Kanfi
 5,642,496 A 6/1997 Kanfi
 6,418,478 B1 7/2002 Ignatius et al.
 6,526,418 B1 2/2003 Midgley et al.
 6,542,972 B2 4/2003 Ignatius et al.
 6,625,623 B1 9/2003 Midgley et al.
 6,658,436 B2 12/2003 Oshinsky et al.
 6,721,767 B2 4/2004 De Meno et al.
 6,760,723 B2 7/2004 Oshinsky et al.
 7,003,641 B2 2/2006 Prahlad et al.
 7,035,880 B1 4/2006 Crescenti et al.
 7,096,330 B1 8/2006 Root et al.
 7,130,970 B2 10/2006 Devassy et al.
 7,162,496 B2 1/2007 Amarendran et al.
 7,174,433 B2 2/2007 Kottomtharayil et al.
 7,246,207 B2 7/2007 Kottomtharayil et al.
 7,315,923 B2 1/2008 Vijayan Retnamma et al.
 7,343,453 B2 3/2008 Prahlad et al.
 7,389,311 B1 6/2008 Crescenti et al.
 7,395,282 B1 7/2008 Crescenti et al.
 7,421,411 B2 9/2008 Kontio et al.
 7,440,982 B2 10/2008 Lu et al.
 7,454,569 B2 11/2008 Kavuri et al.
 7,483,925 B2 1/2009 Koskimies et al.
 7,490,207 B2 2/2009 Amarendran et al.
 7,500,053 B1 3/2009 Kavuri et al.
 7,523,276 B1 4/2009 Shankar
 7,529,782 B2 5/2009 Prahlad et al.
 7,536,291 B1 5/2009 Vijayan Retnamma et al.
 7,543,125 B2 6/2009 Gokhale
 7,546,324 B2 6/2009 Prahlad et al.
 7,603,386 B2 10/2009 Amarendran et al.
 7,606,844 B2 10/2009 Kottomtharayil
 7,613,752 B2 11/2009 Prahlad et al.
 7,617,253 B2 11/2009 Prahlad et al.
 7,617,262 B2 11/2009 Prahlad et al.
 7,620,710 B2 11/2009 Kottomtharayil et al.
 7,636,743 B2 12/2009 Erofeev
 7,651,593 B2 1/2010 Prahlad et al.
 7,657,550 B2 2/2010 Prahlad et al.
 7,660,807 B2 2/2010 Prahlad et al.
 7,661,028 B2 2/2010 Erofeev
 7,747,579 B2 6/2010 Prahlad et al.
 7,801,864 B2 9/2010 Prahlad et al.

7,809,914 B2 10/2010 Kottomtharayil et al.
 7,814,053 B1 10/2010 Pivowar et al.
 8,069,218 B1 11/2011 Tormasov et al.
 8,170,995 B2 5/2012 Prahlad et al.
 8,229,888 B1 7/2012 Roskind et al.
 8,229,954 B2 7/2012 Kottomtharayil et al.
 8,285,681 B2 10/2012 Prahlad et al.
 8,307,177 B2 11/2012 Prahlad et al.
 8,364,652 B2 1/2013 Vijayan Retnamma et al.
 8,370,542 B2 2/2013 Lu et al.
 8,554,951 B2 10/2013 Barton et al.
 8,578,120 B2 11/2013 Attarde et al.
 8,630,980 B2 1/2014 Piwonka et al.
 8,655,840 B2 2/2014 Simelius
 9,262,435 B2 2/2016 Vibhor et al.
 9,317,223 B2 4/2016 Reohr et al.
 9,336,226 B2 5/2016 Vibhor et al.
 9,367,549 B2 6/2016 Hrebicek et al.
 9,430,491 B2 8/2016 Vibhor et al.
 9,483,201 B2 11/2016 Kripalani et al.
 9,483,489 B2 11/2016 Varadharajan et al.
 9,495,251 B2 11/2016 Kottomtharayil
 9,588,849 B2 3/2017 Sinha
 9,753,816 B2 9/2017 Mehta et al.
 9,898,481 B2 2/2018 Vibhor et al.
 9,928,005 B2 3/2018 Sinha
 9,952,934 B2 4/2018 Sinha
 2002/0120791 A1* 8/2002 Somalwar G06F 11/1435
 719/330
 2003/0120685 A1 6/2003 Duncombe et al.
 2004/0064488 A1 4/2004 Sinha
 2005/0071386 A1 3/2005 Wolfgang et al.
 2005/0144195 A1 6/2005 Hesselink et al.
 2006/0224846 A1 10/2006 Amarendran et al.
 2007/0143371 A1 6/2007 Kottomtharayil
 2007/0174569 A1 7/2007 Schnapp
 2008/0209146 A1 8/2008 Imazu
 2009/0276698 A1 11/2009 Clarke et al.
 2009/0319534 A1 12/2009 Gokhale
 2009/0319540 A1 12/2009 Suthar et al.
 2010/0076933 A1 3/2010 Hamilton et al.
 2010/0138387 A1 6/2010 Simelius
 2010/0180094 A1* 7/2010 Min G06F 11/2058
 711/162
 2011/0196957 A1 8/2011 Ayachitula
 2012/0131684 A1 5/2012 Lynch
 2012/0150818 A1 6/2012 Vijayan Retnamma et al.
 2012/0150826 A1 6/2012 Vijayan Retnamma et al.
 2012/0296870 A1 11/2012 Fu et al.
 2013/0238696 A1 9/2013 Cotelo
 2013/0262800 A1 10/2013 Goodman
 2014/0053227 A1 2/2014 Ruppin et al.
 2014/0143201 A1 5/2014 Hwang
 2014/0172793 A1 6/2014 Stritzel et al.
 2014/0172798 A1 6/2014 Falkenburg et al.
 2014/0173229 A1 6/2014 Reohr et al.
 2014/0173232 A1 6/2014 Reohr et al.
 2014/0188804 A1 7/2014 Gokhale et al.
 2014/0201140 A1 7/2014 Vibhor et al.
 2014/0214764 A1 7/2014 Robbin et al.
 2014/0281317 A1 9/2014 Garman
 2014/0289191 A1 9/2014 Chan et al.
 2015/0193312 A1 7/2015 Nanivadekar et al.
 2016/0127452 A1 5/2016 Newman et al.
 2016/0162369 A1 6/2016 Ahn et al.
 2016/0210195 A1 7/2016 Sinha
 2016/0246815 A1 8/2016 Vibhor et al.
 2017/0090795 A1 3/2017 Varadharajan et al.
 2017/0109370 A1 4/2017 Newhouse
 2017/0131912 A1 5/2017 Sinha
 2017/0235756 A1 8/2017 Mehta et al.

FOREIGN PATENT DOCUMENTS

EP 0467546 1/1992
 EP 0541281 5/1993
 EP 0774715 5/1997
 EP 0809184 11/1997
 EP 0899662 3/1999

(56)

References Cited

FOREIGN PATENT DOCUMENTS

EP	0981090	2/2000
WO	WO 1995/13580	5/1995
WO	WO 1999/12098	3/1999
WO	WO 2006/052872	5/2006

OTHER PUBLICATIONS

Cabrera, et al. "ADSM: A Multi-Platform, Scalable, Back-up and Archive Mass Storage System," Digest of Papers, Comcon '95, Proceedings of the 40th IEEE Computer Society International Conference, Mar. 5, 1995-Mar. 9, 1995, pp. 420-427, San Francisco, CA.

Eitel, "Backup and Storage Management in Distributed Heterogeneous Environments," IEEE, 1994, pp. 124-126.

Huff, KL, "Data Set Usage Sequence Number," IBM Technical Disclosure Bulletin, vol. 24, No. 5, Oct. 1981 New York, US, pp. 2404-2406.

Rosenblum et al., "The Design and Implementation of a Log-Structure File System," Operating Systems Review SIGOPS, vol. 25, No. 5, New York, US, pp. 1-15 (May 1991).

* cited by examiner

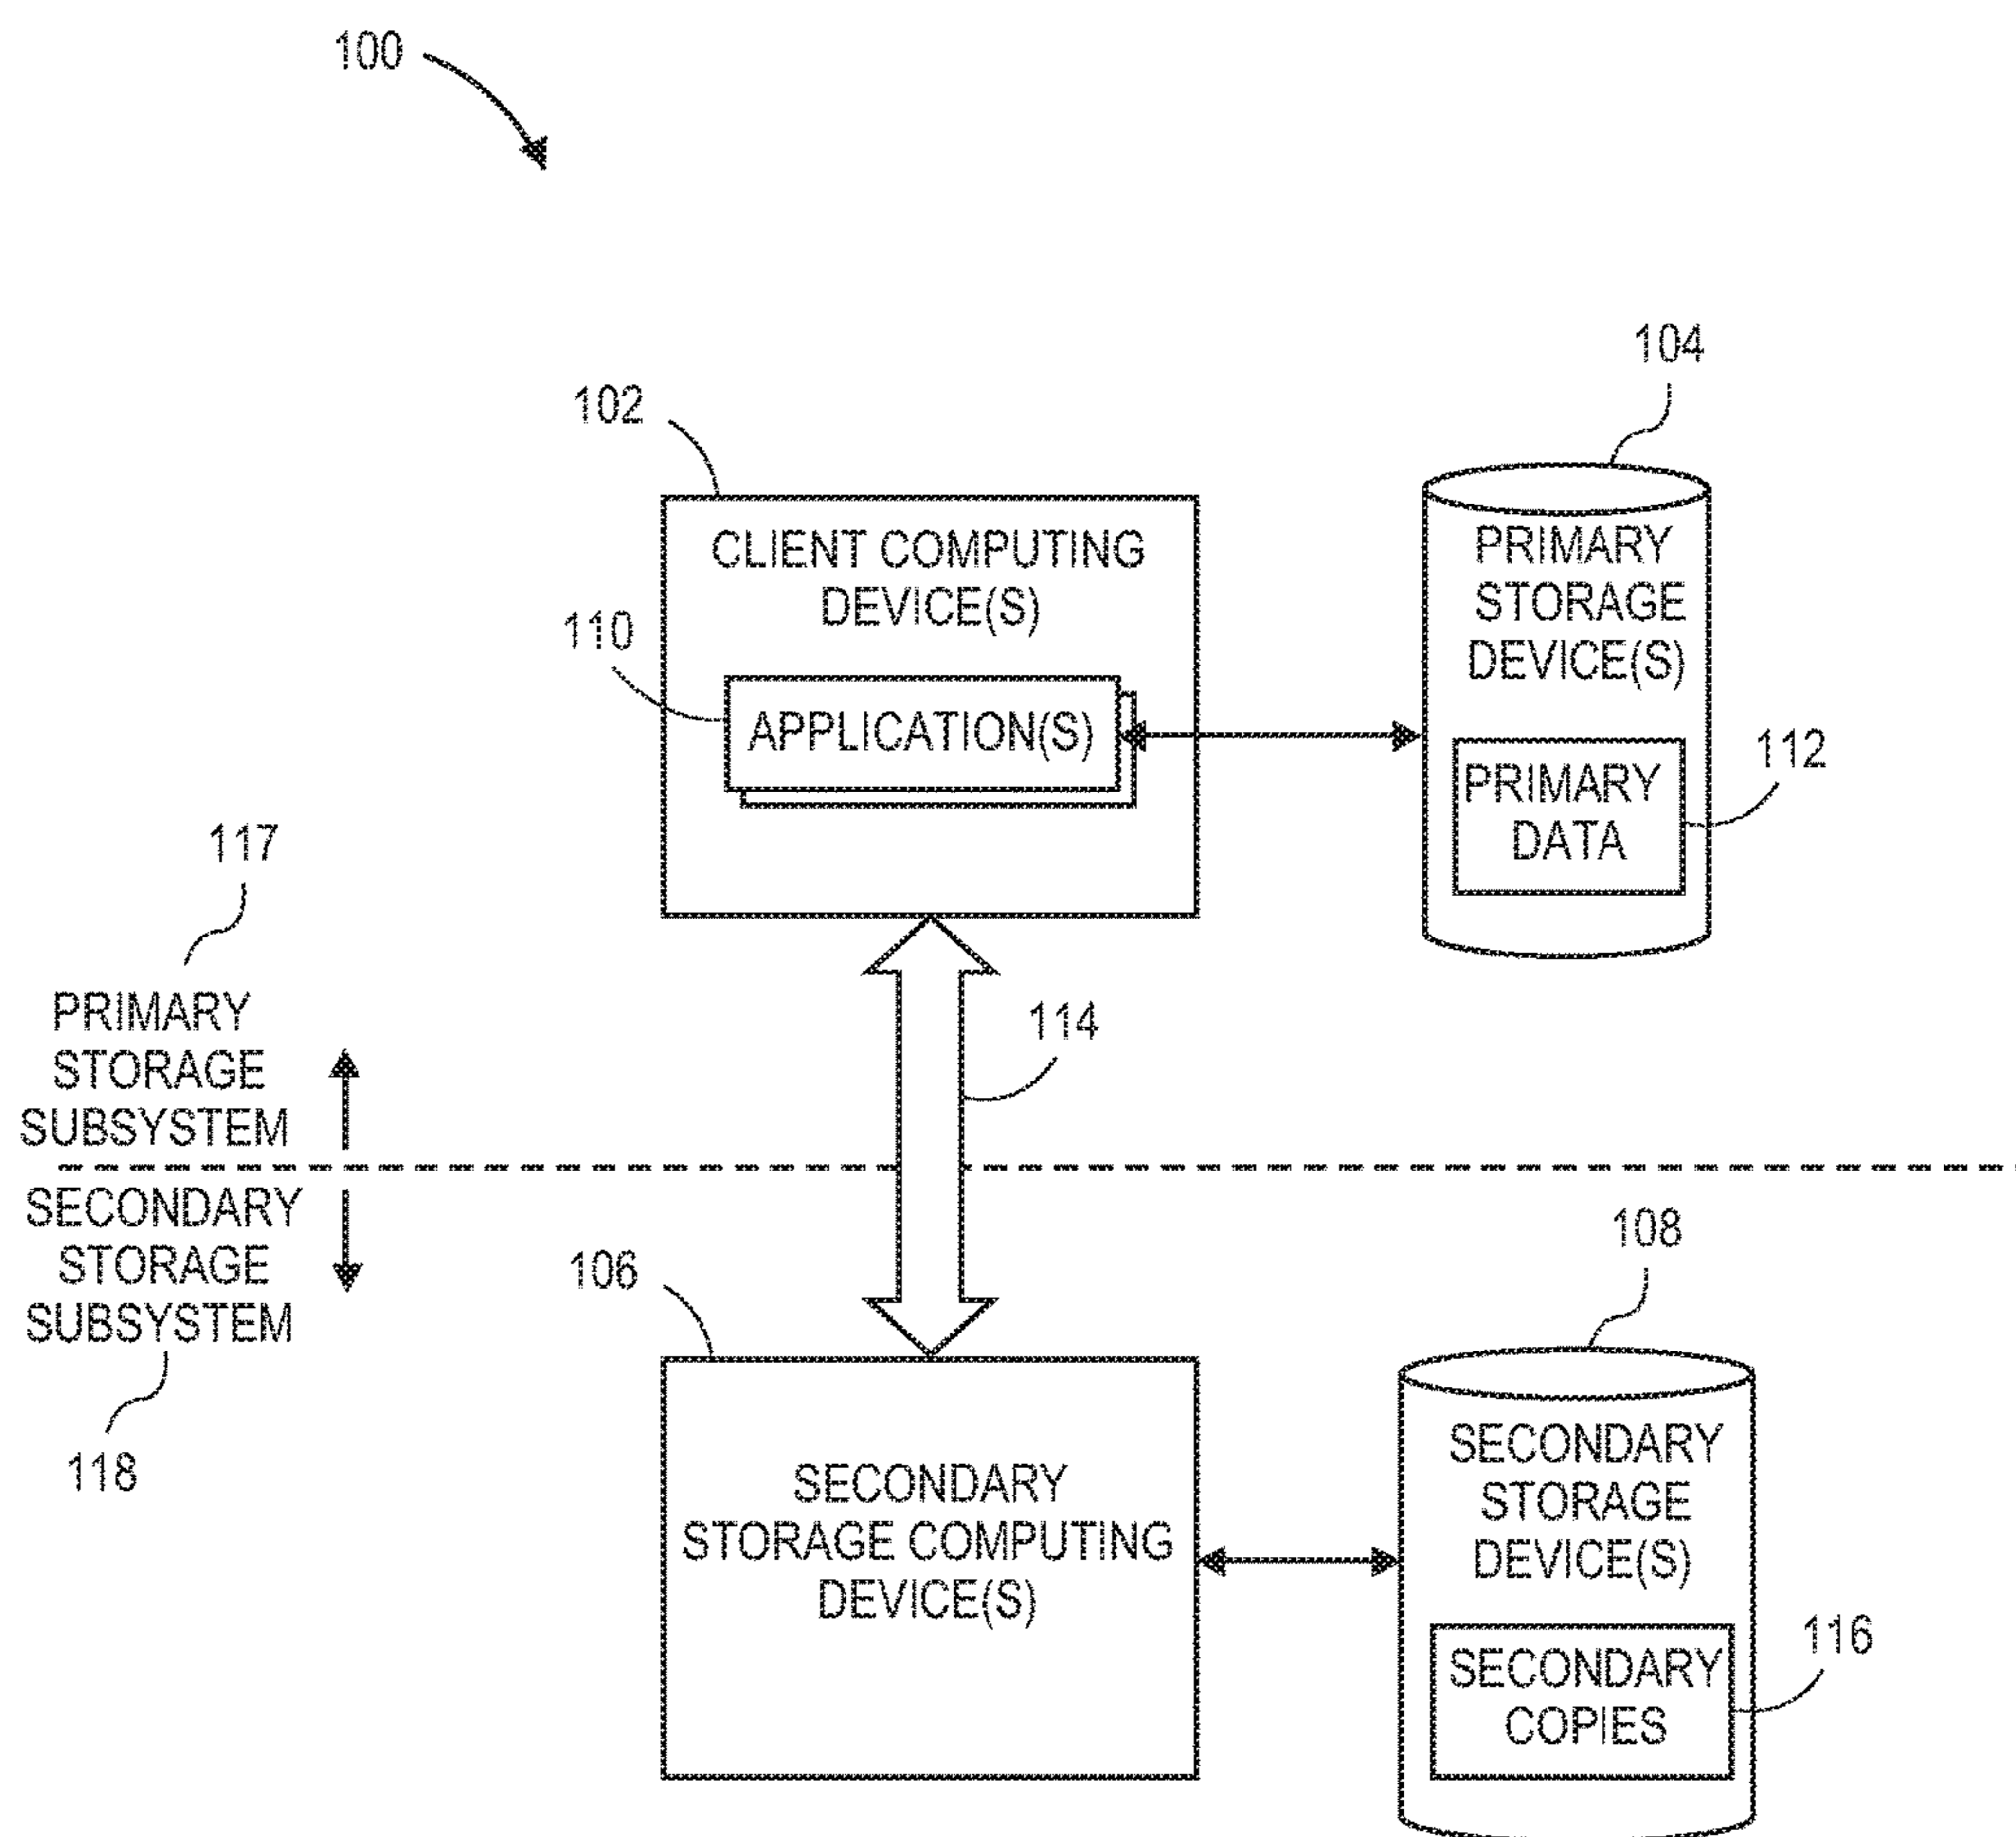


FIG. 1A

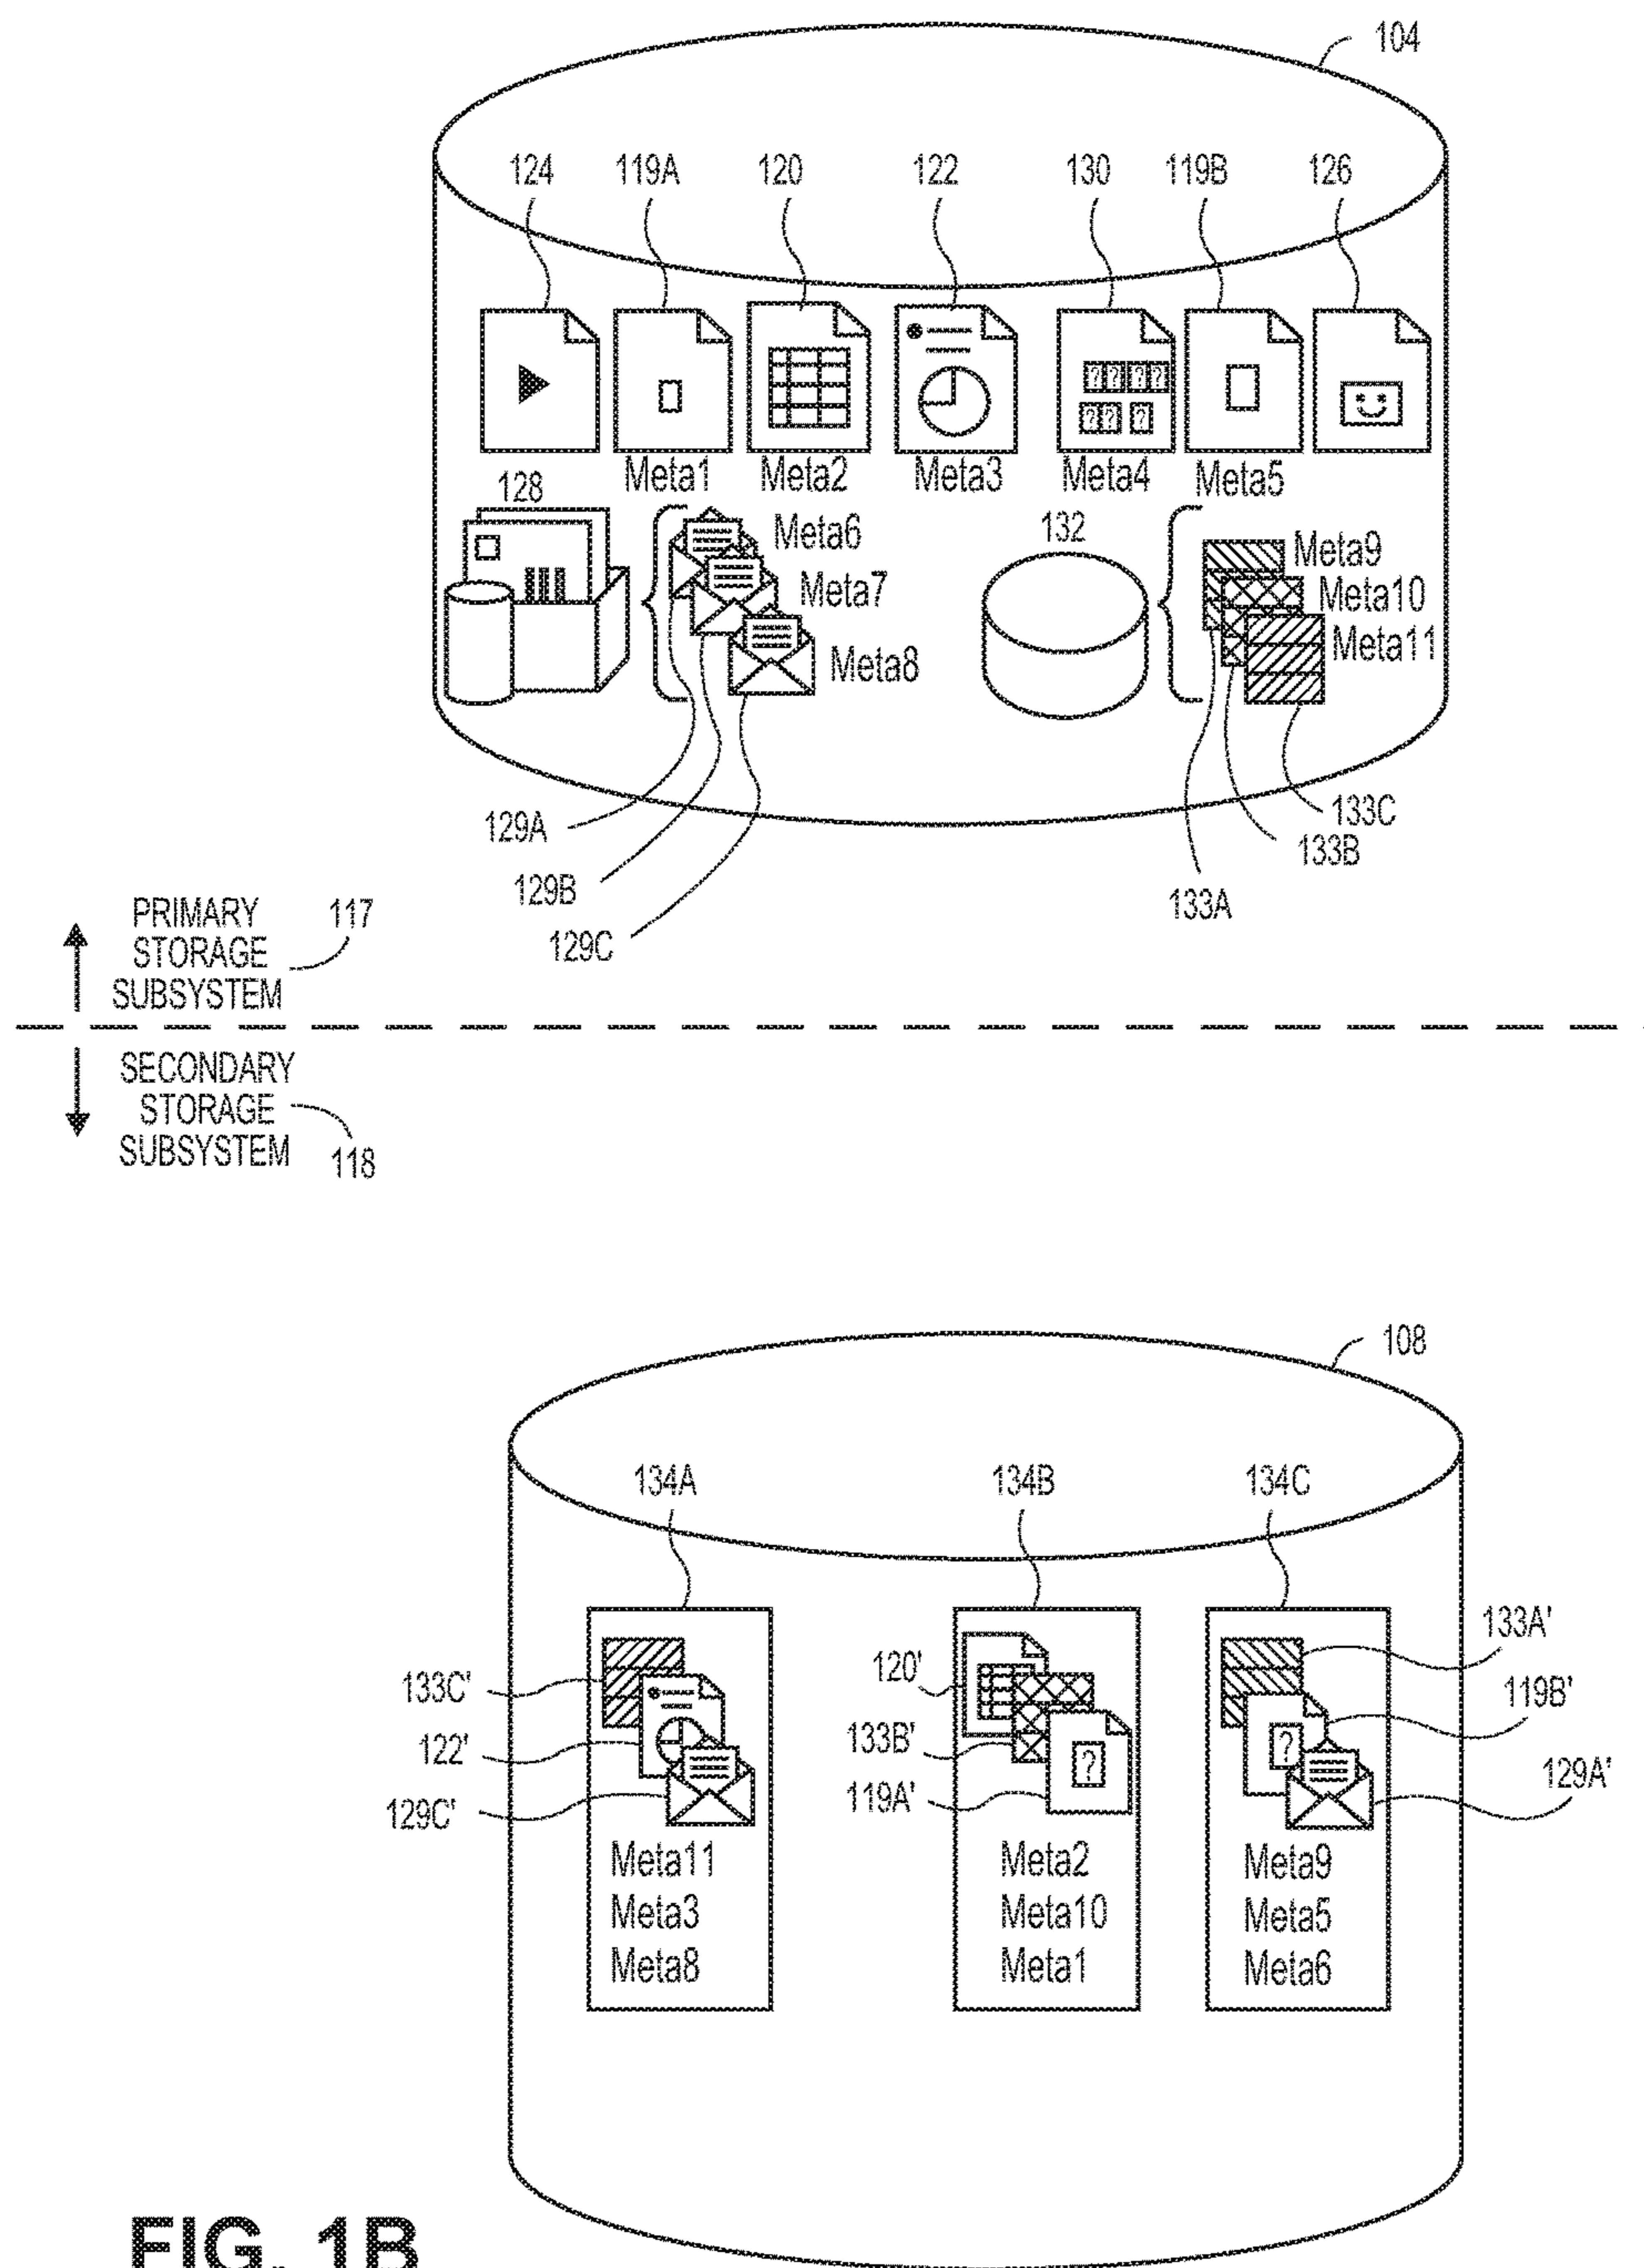


FIG. 1B

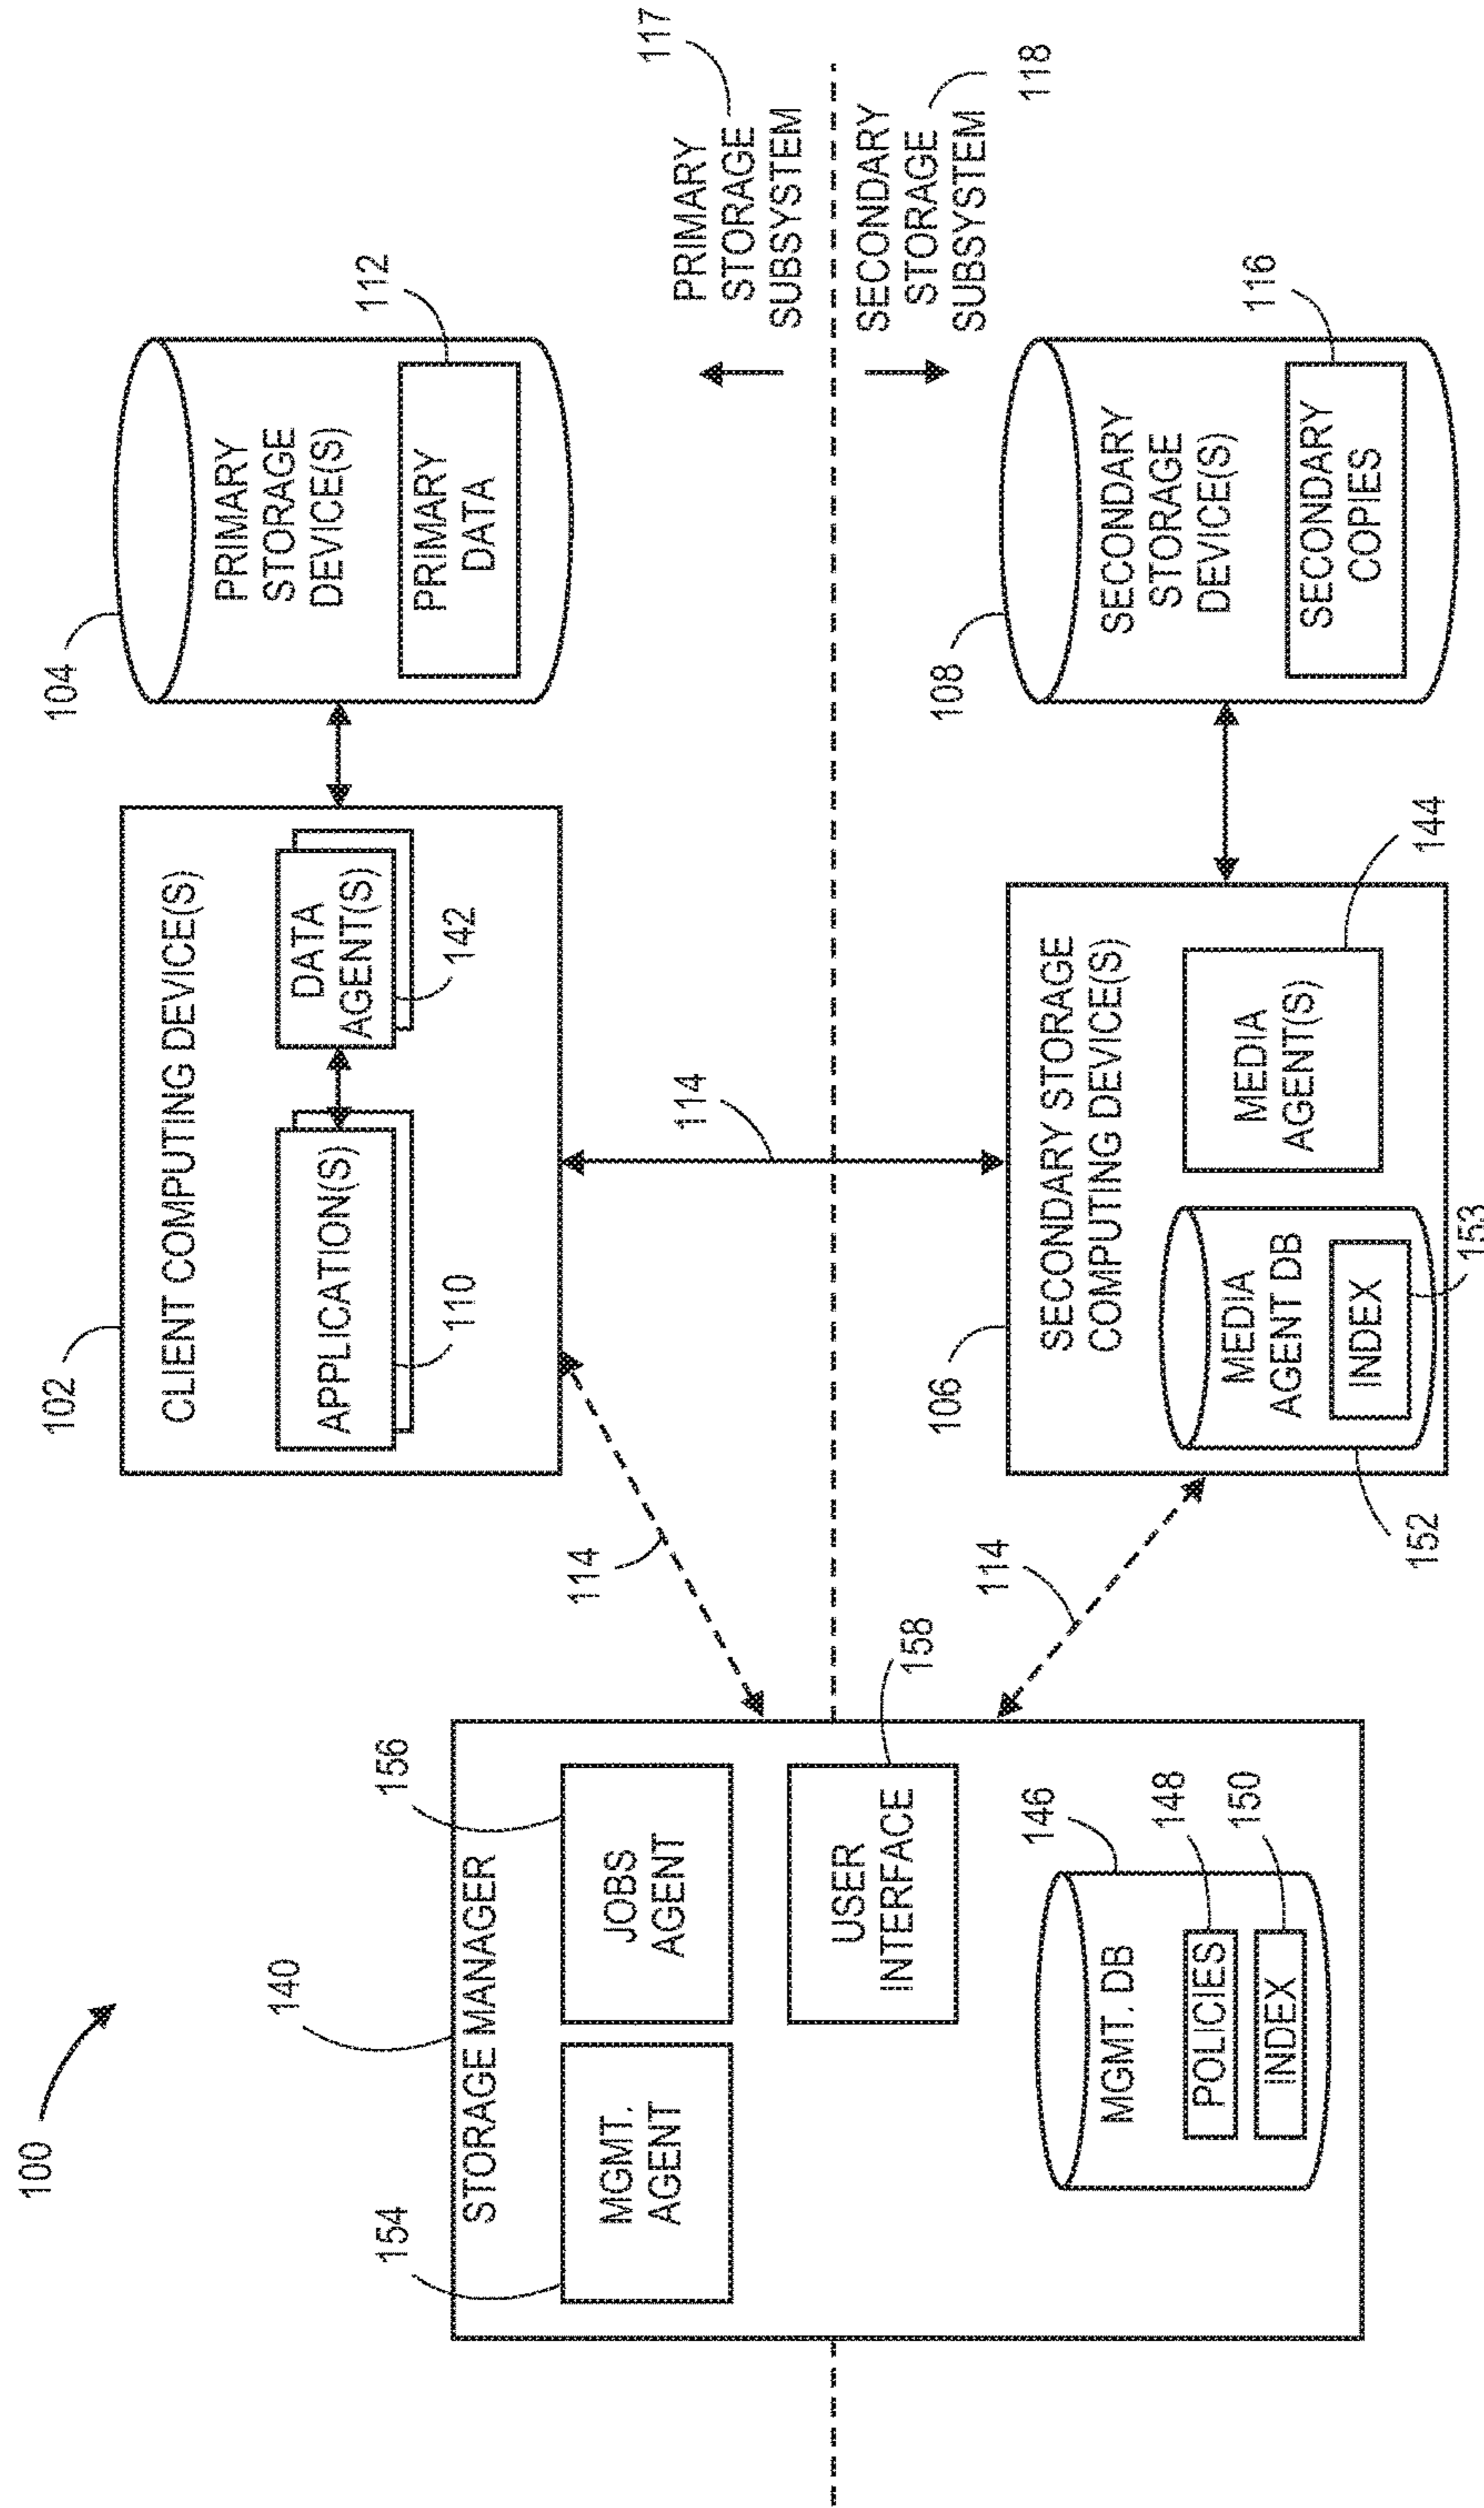


FIG. 1C

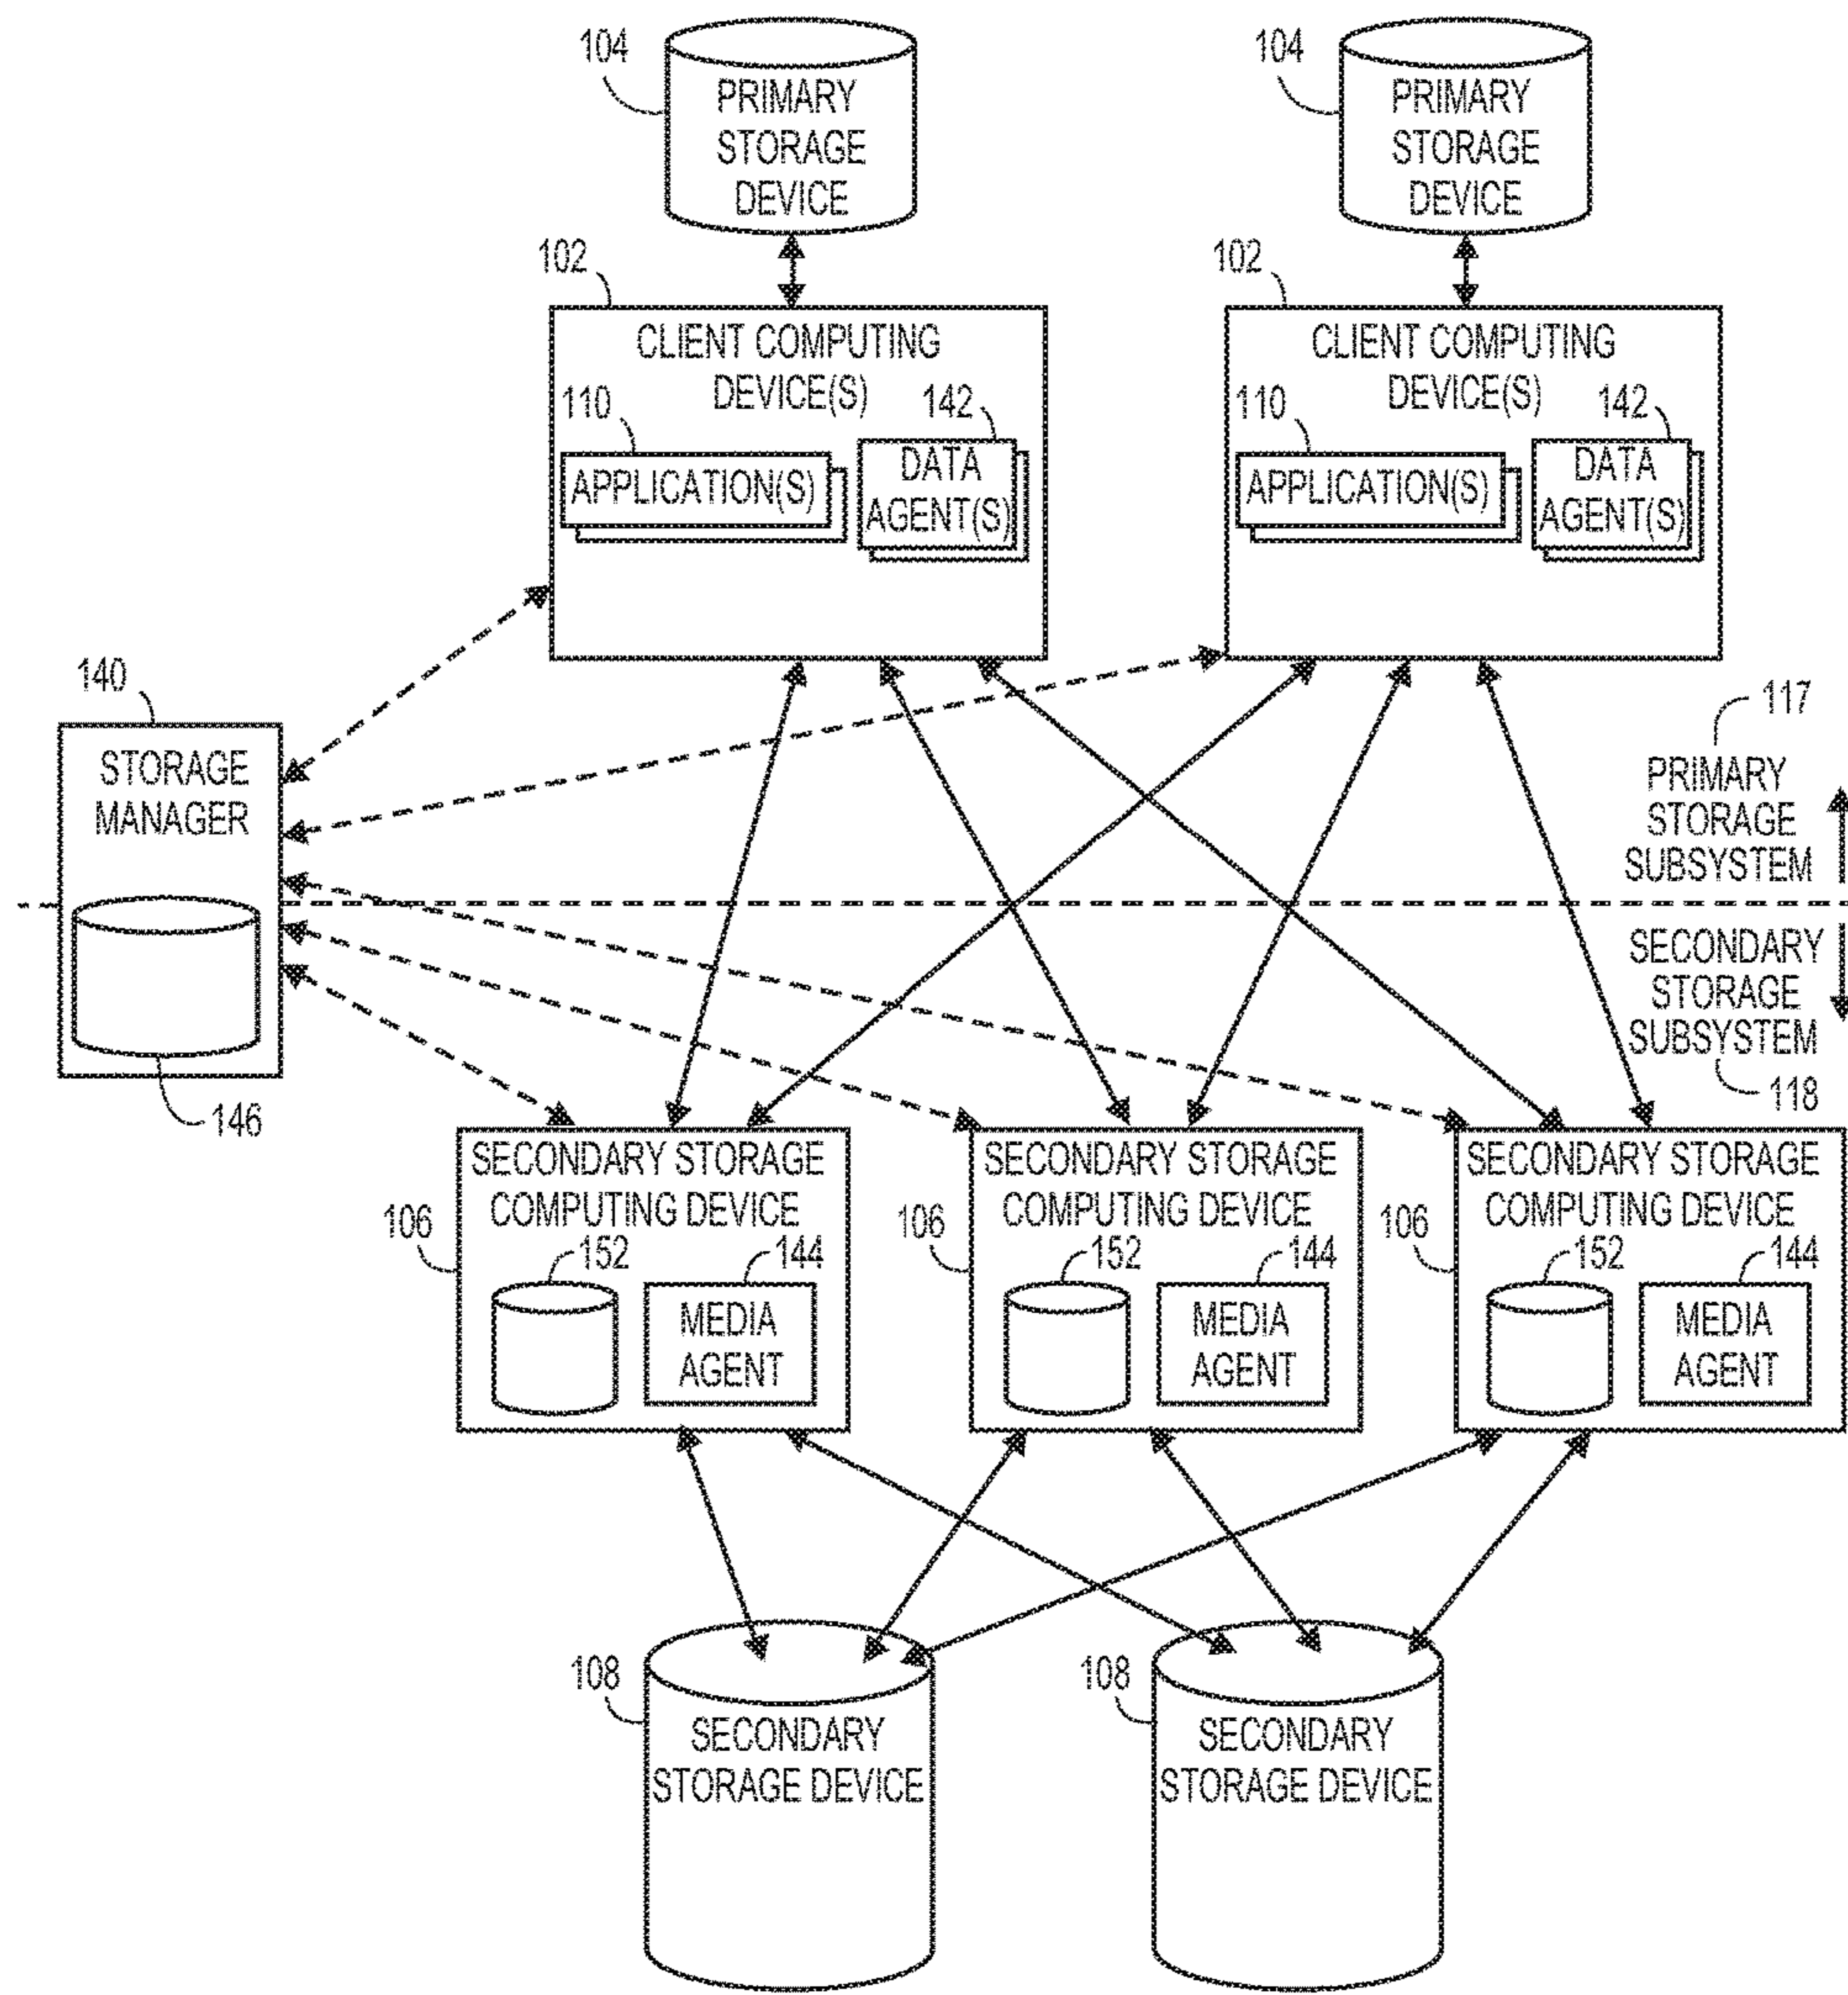


FIG. 1D

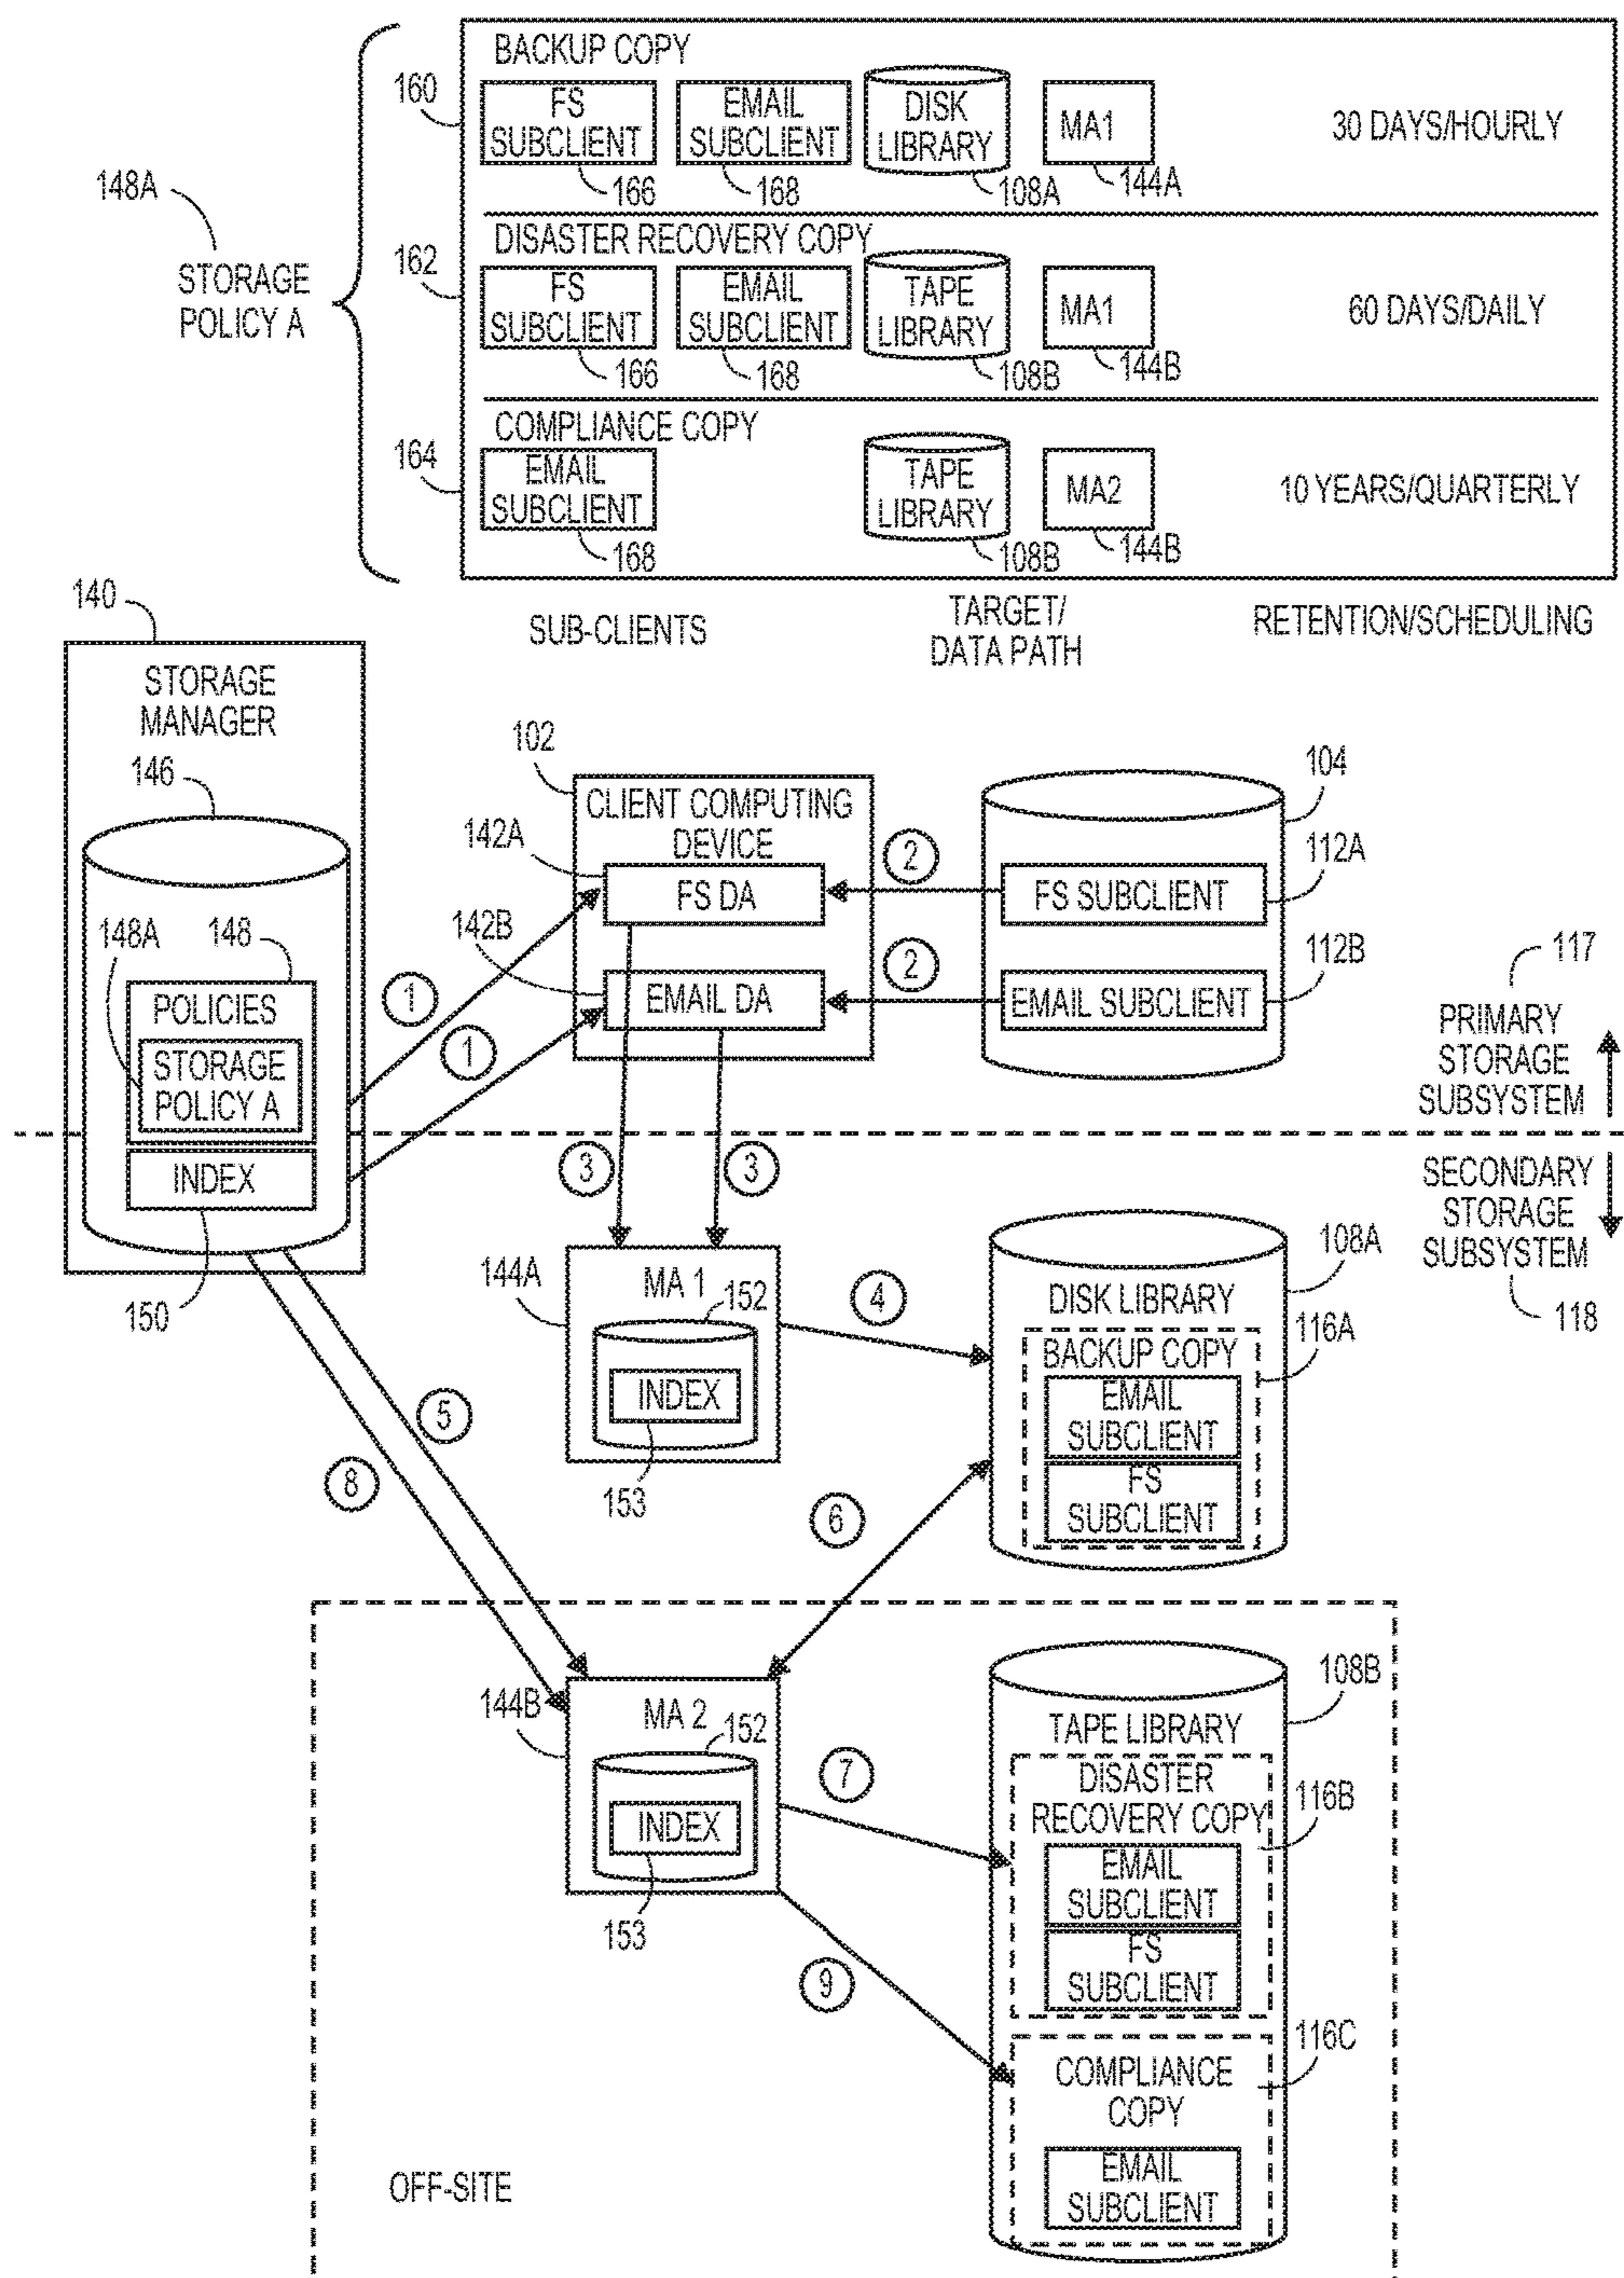


FIG. 1E

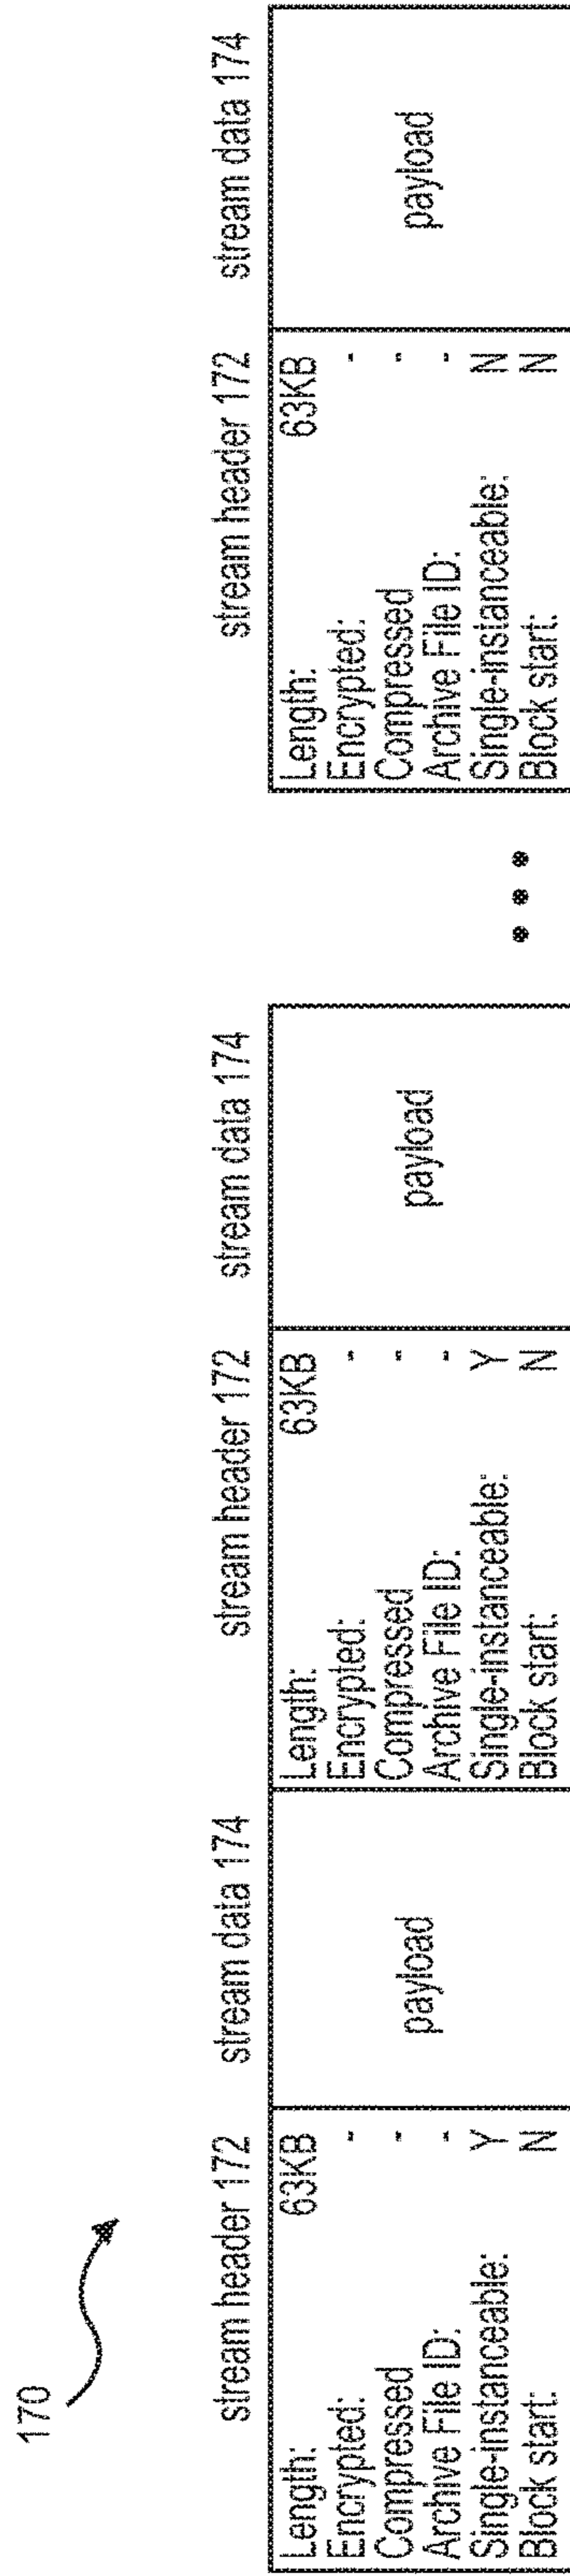


FIG. 1F

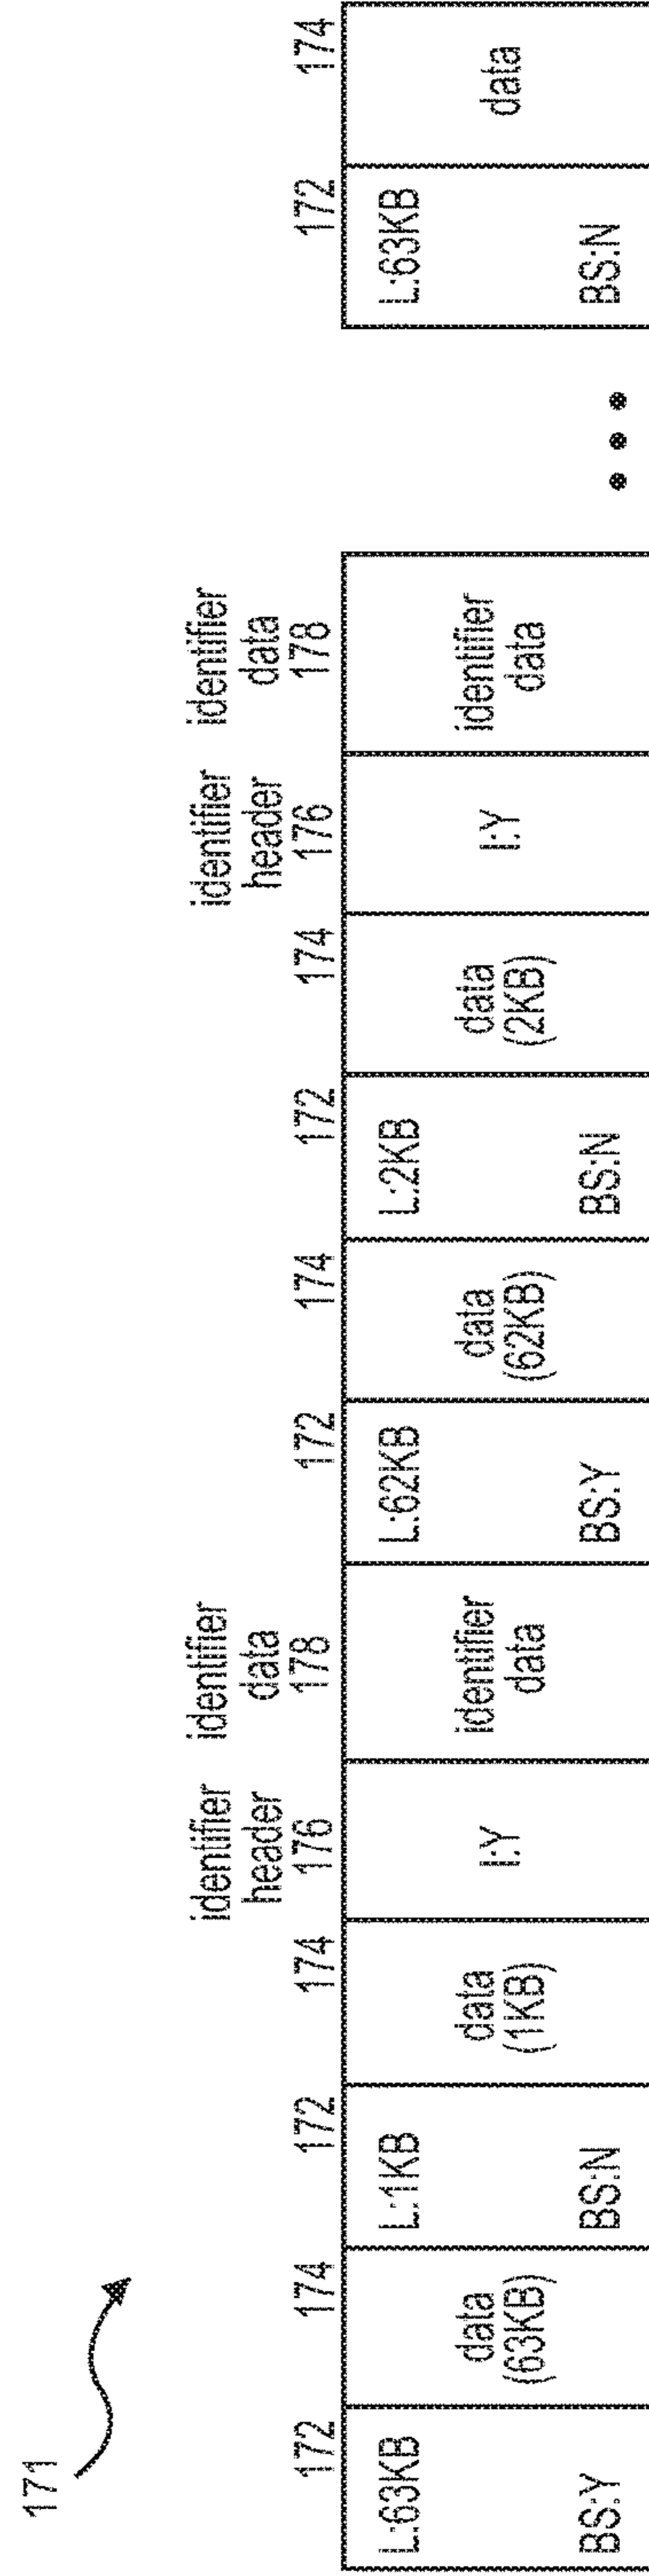


FIG. 1G

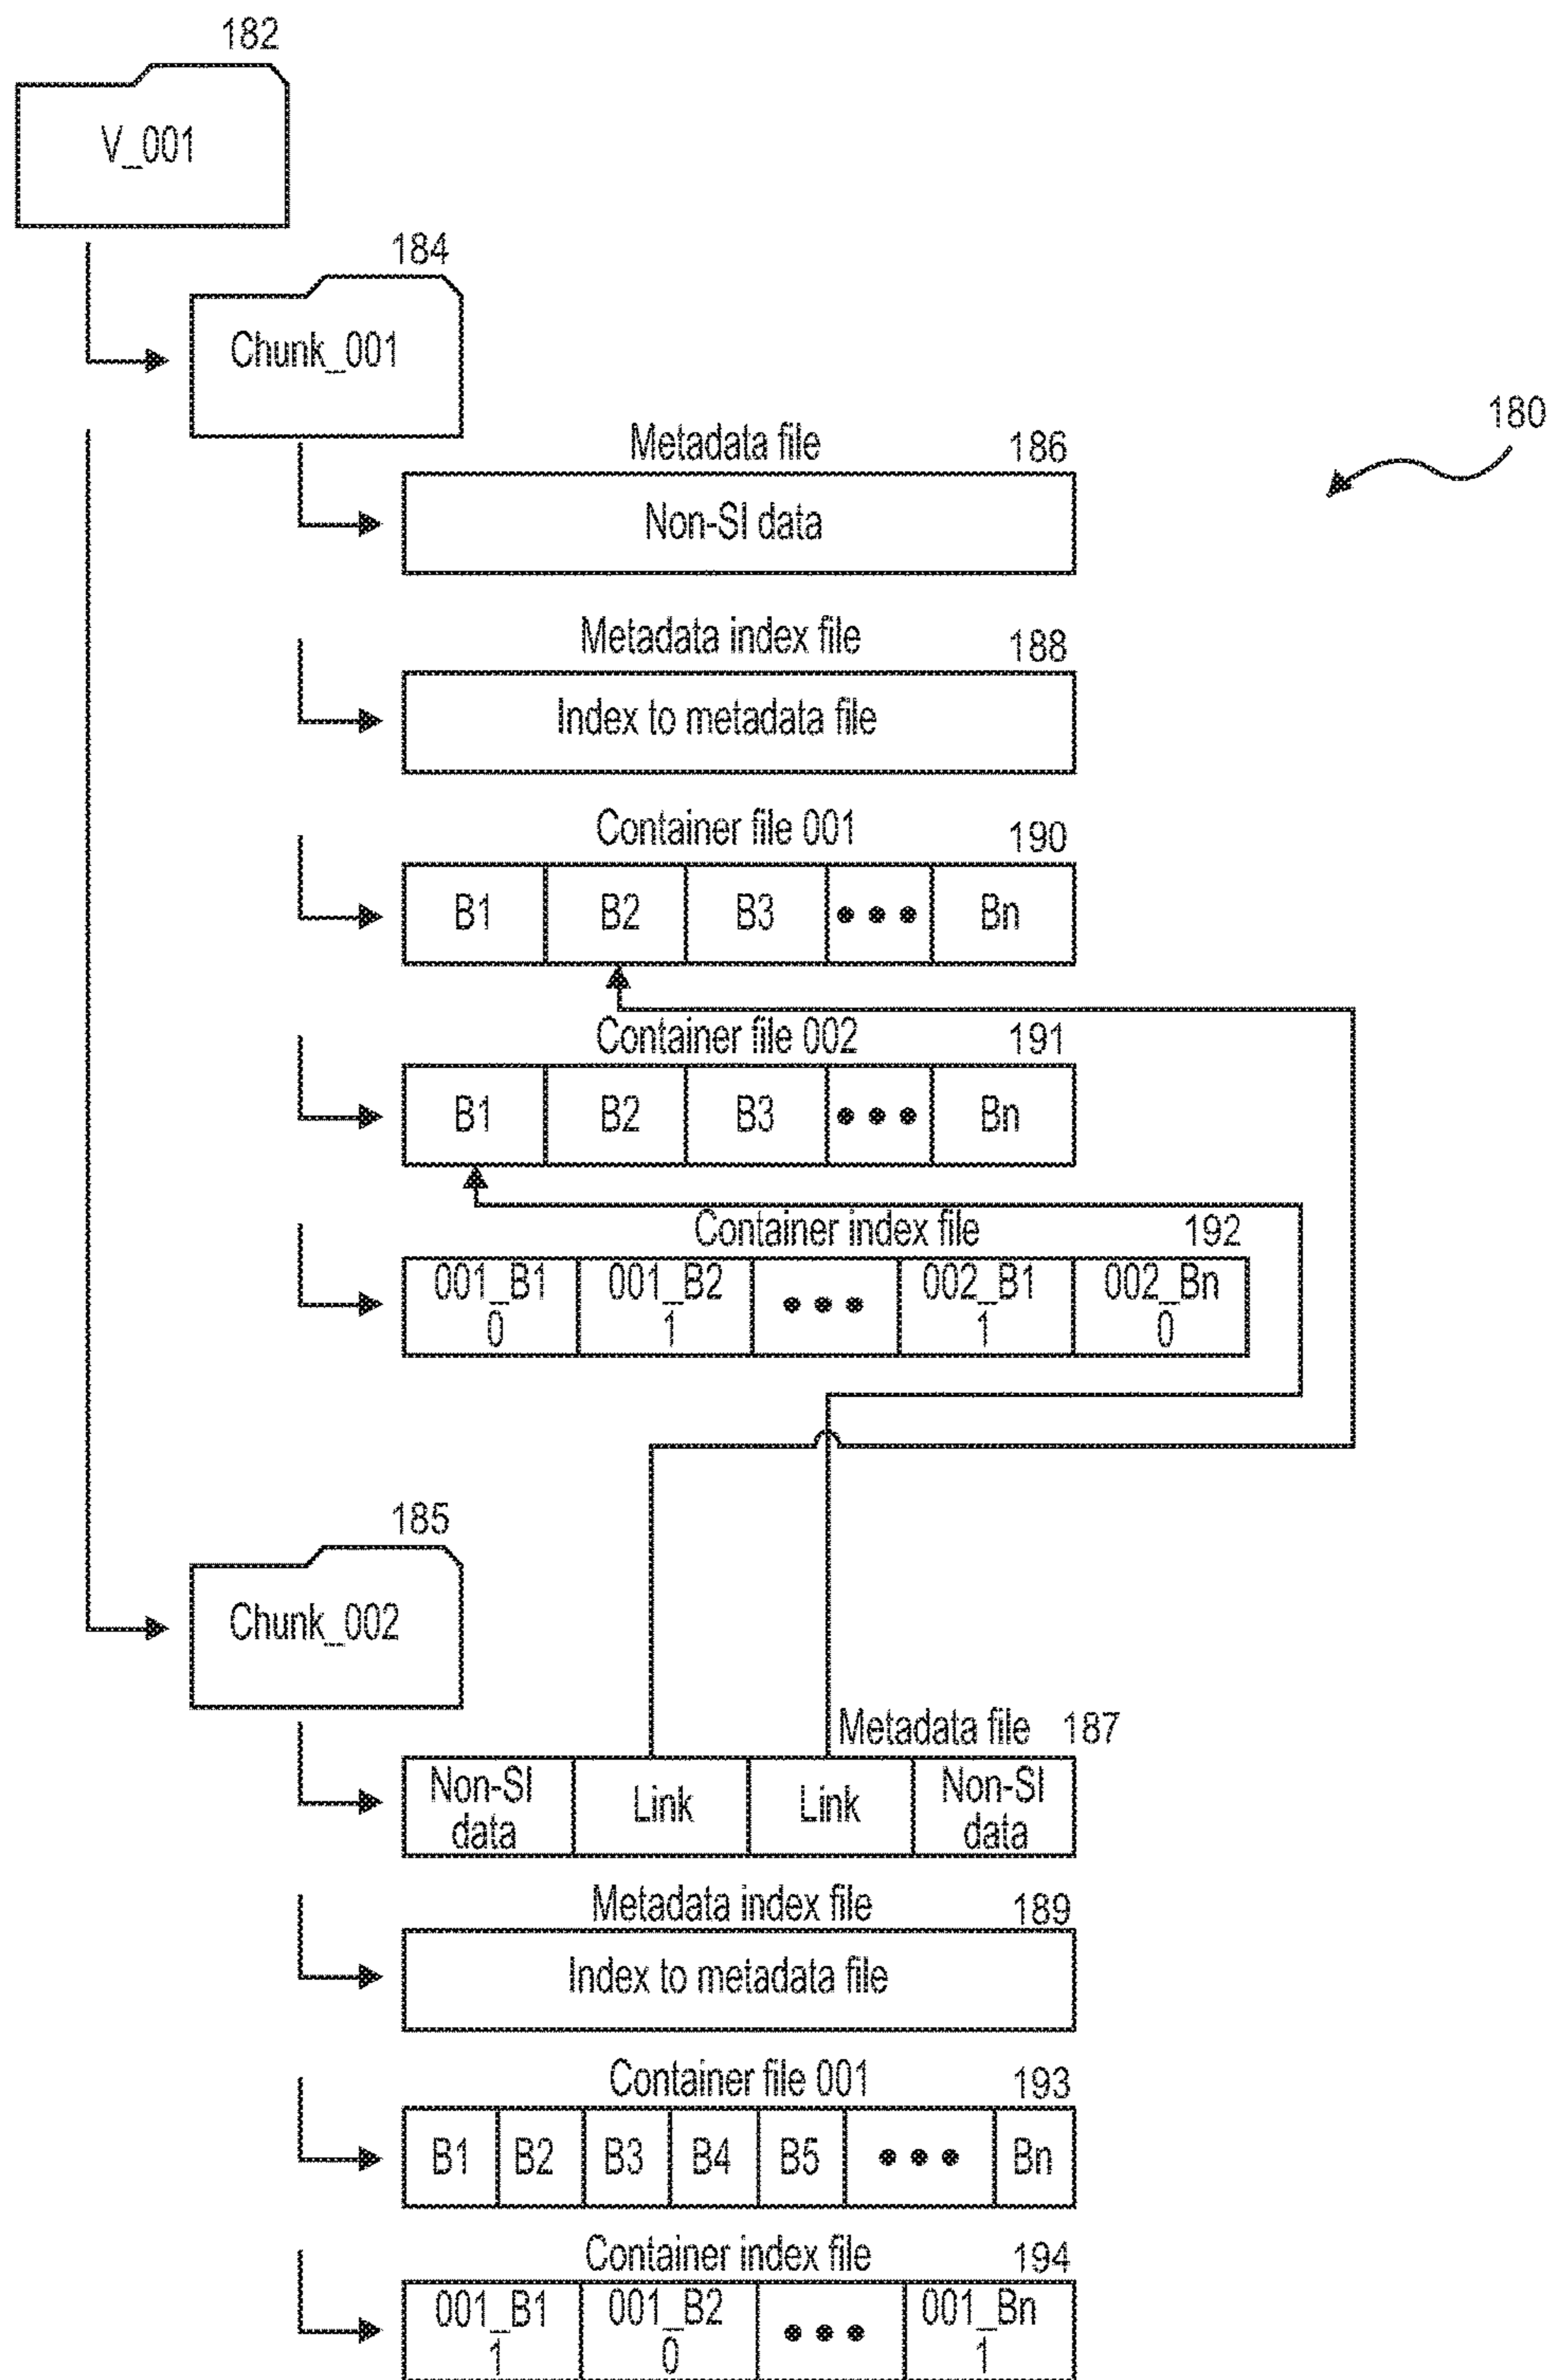


FIG. 1H

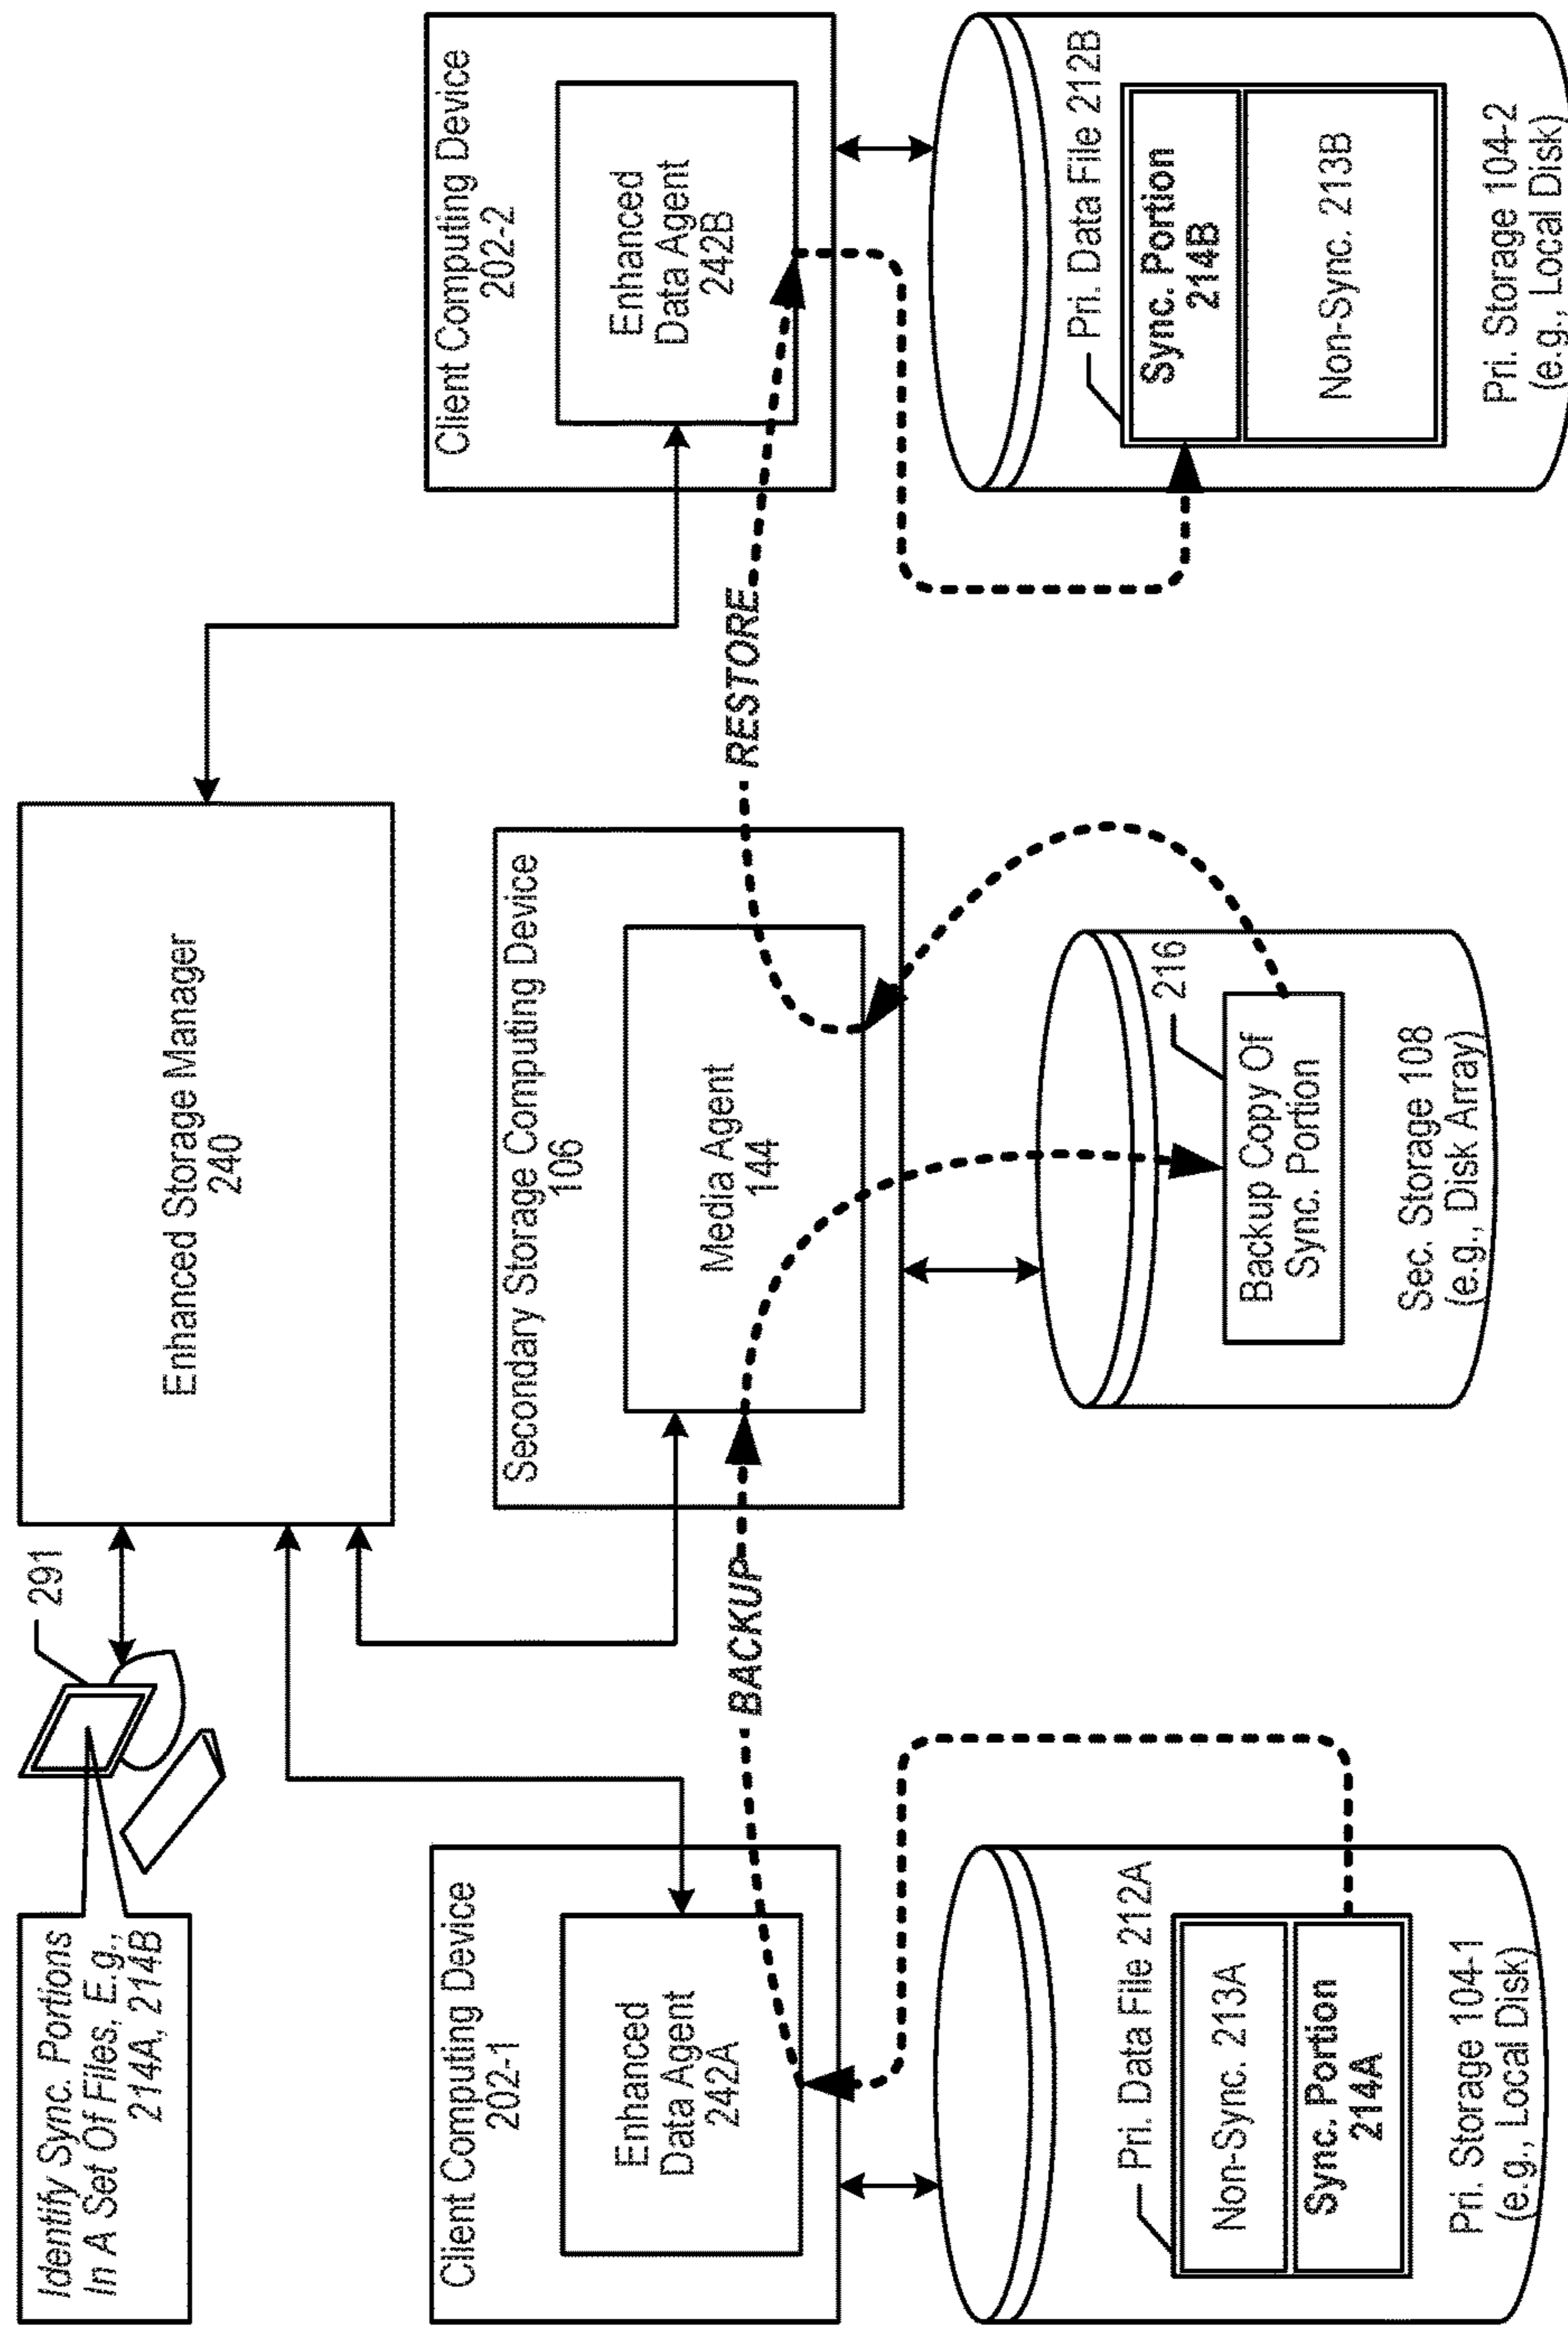


FIG. 2 Storage Mgmt. System 200 For Partially Synchronizing Primary Data Files Based On Synchronizing Portions Thereof

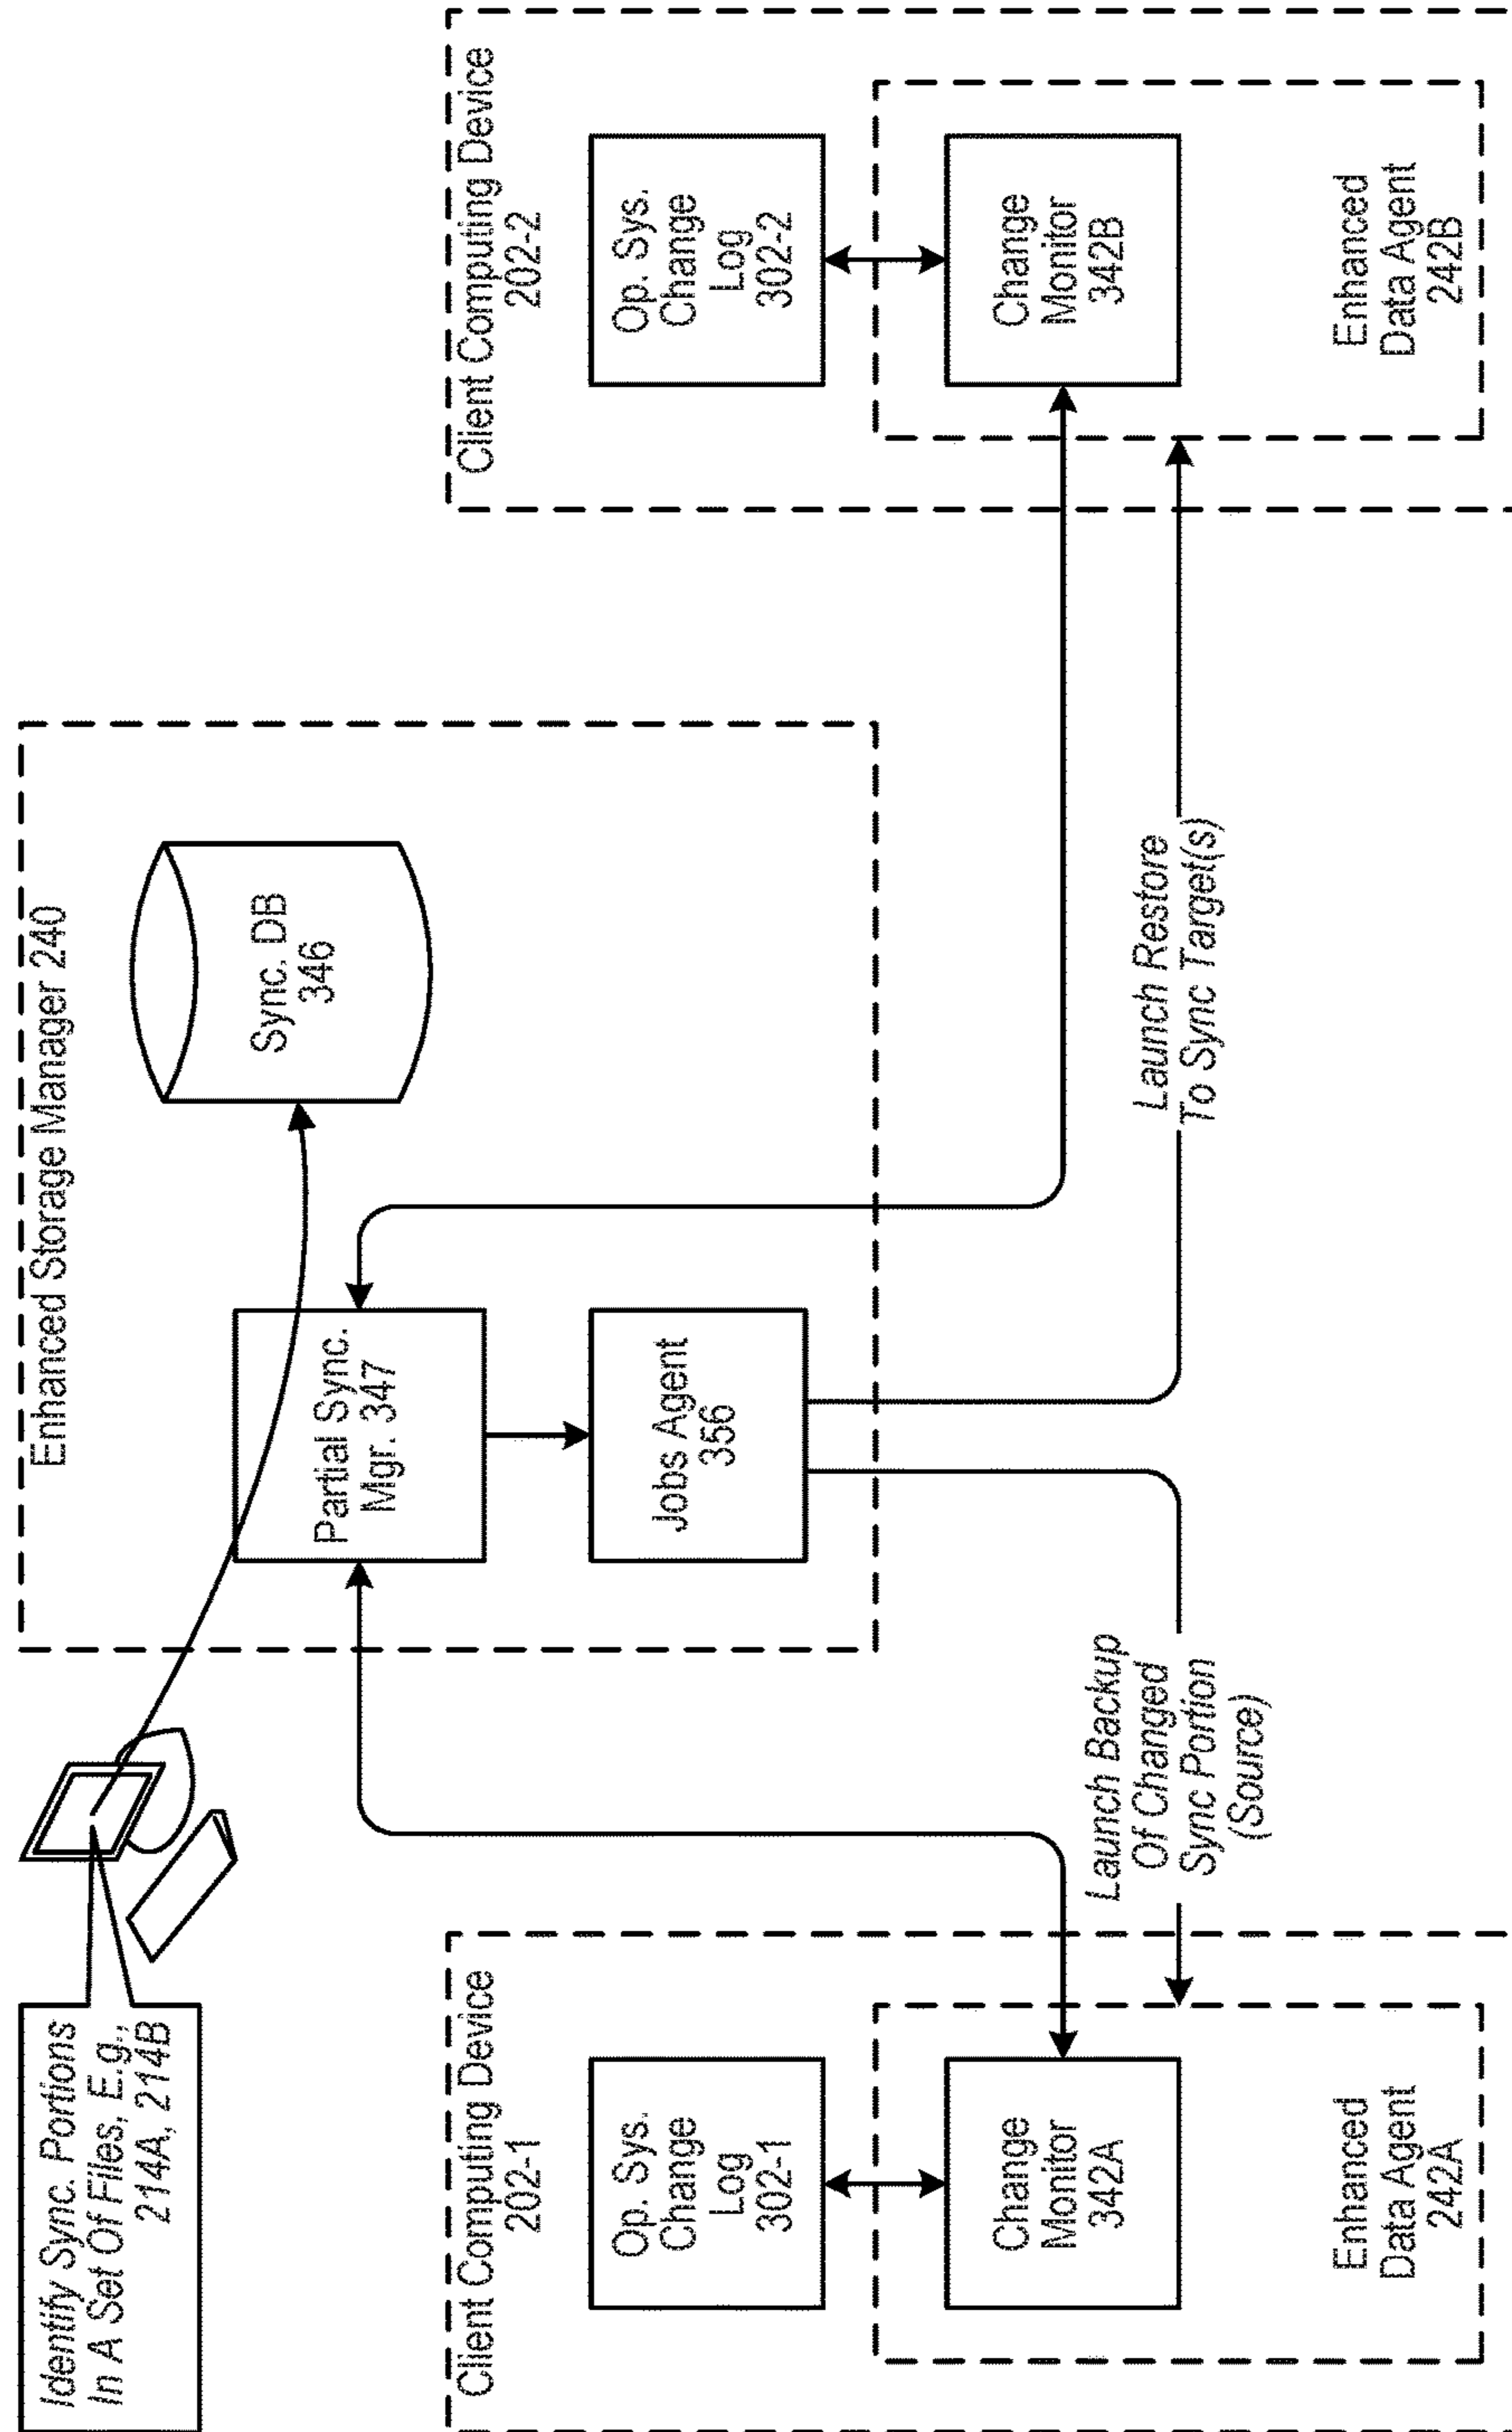


FIG. 3

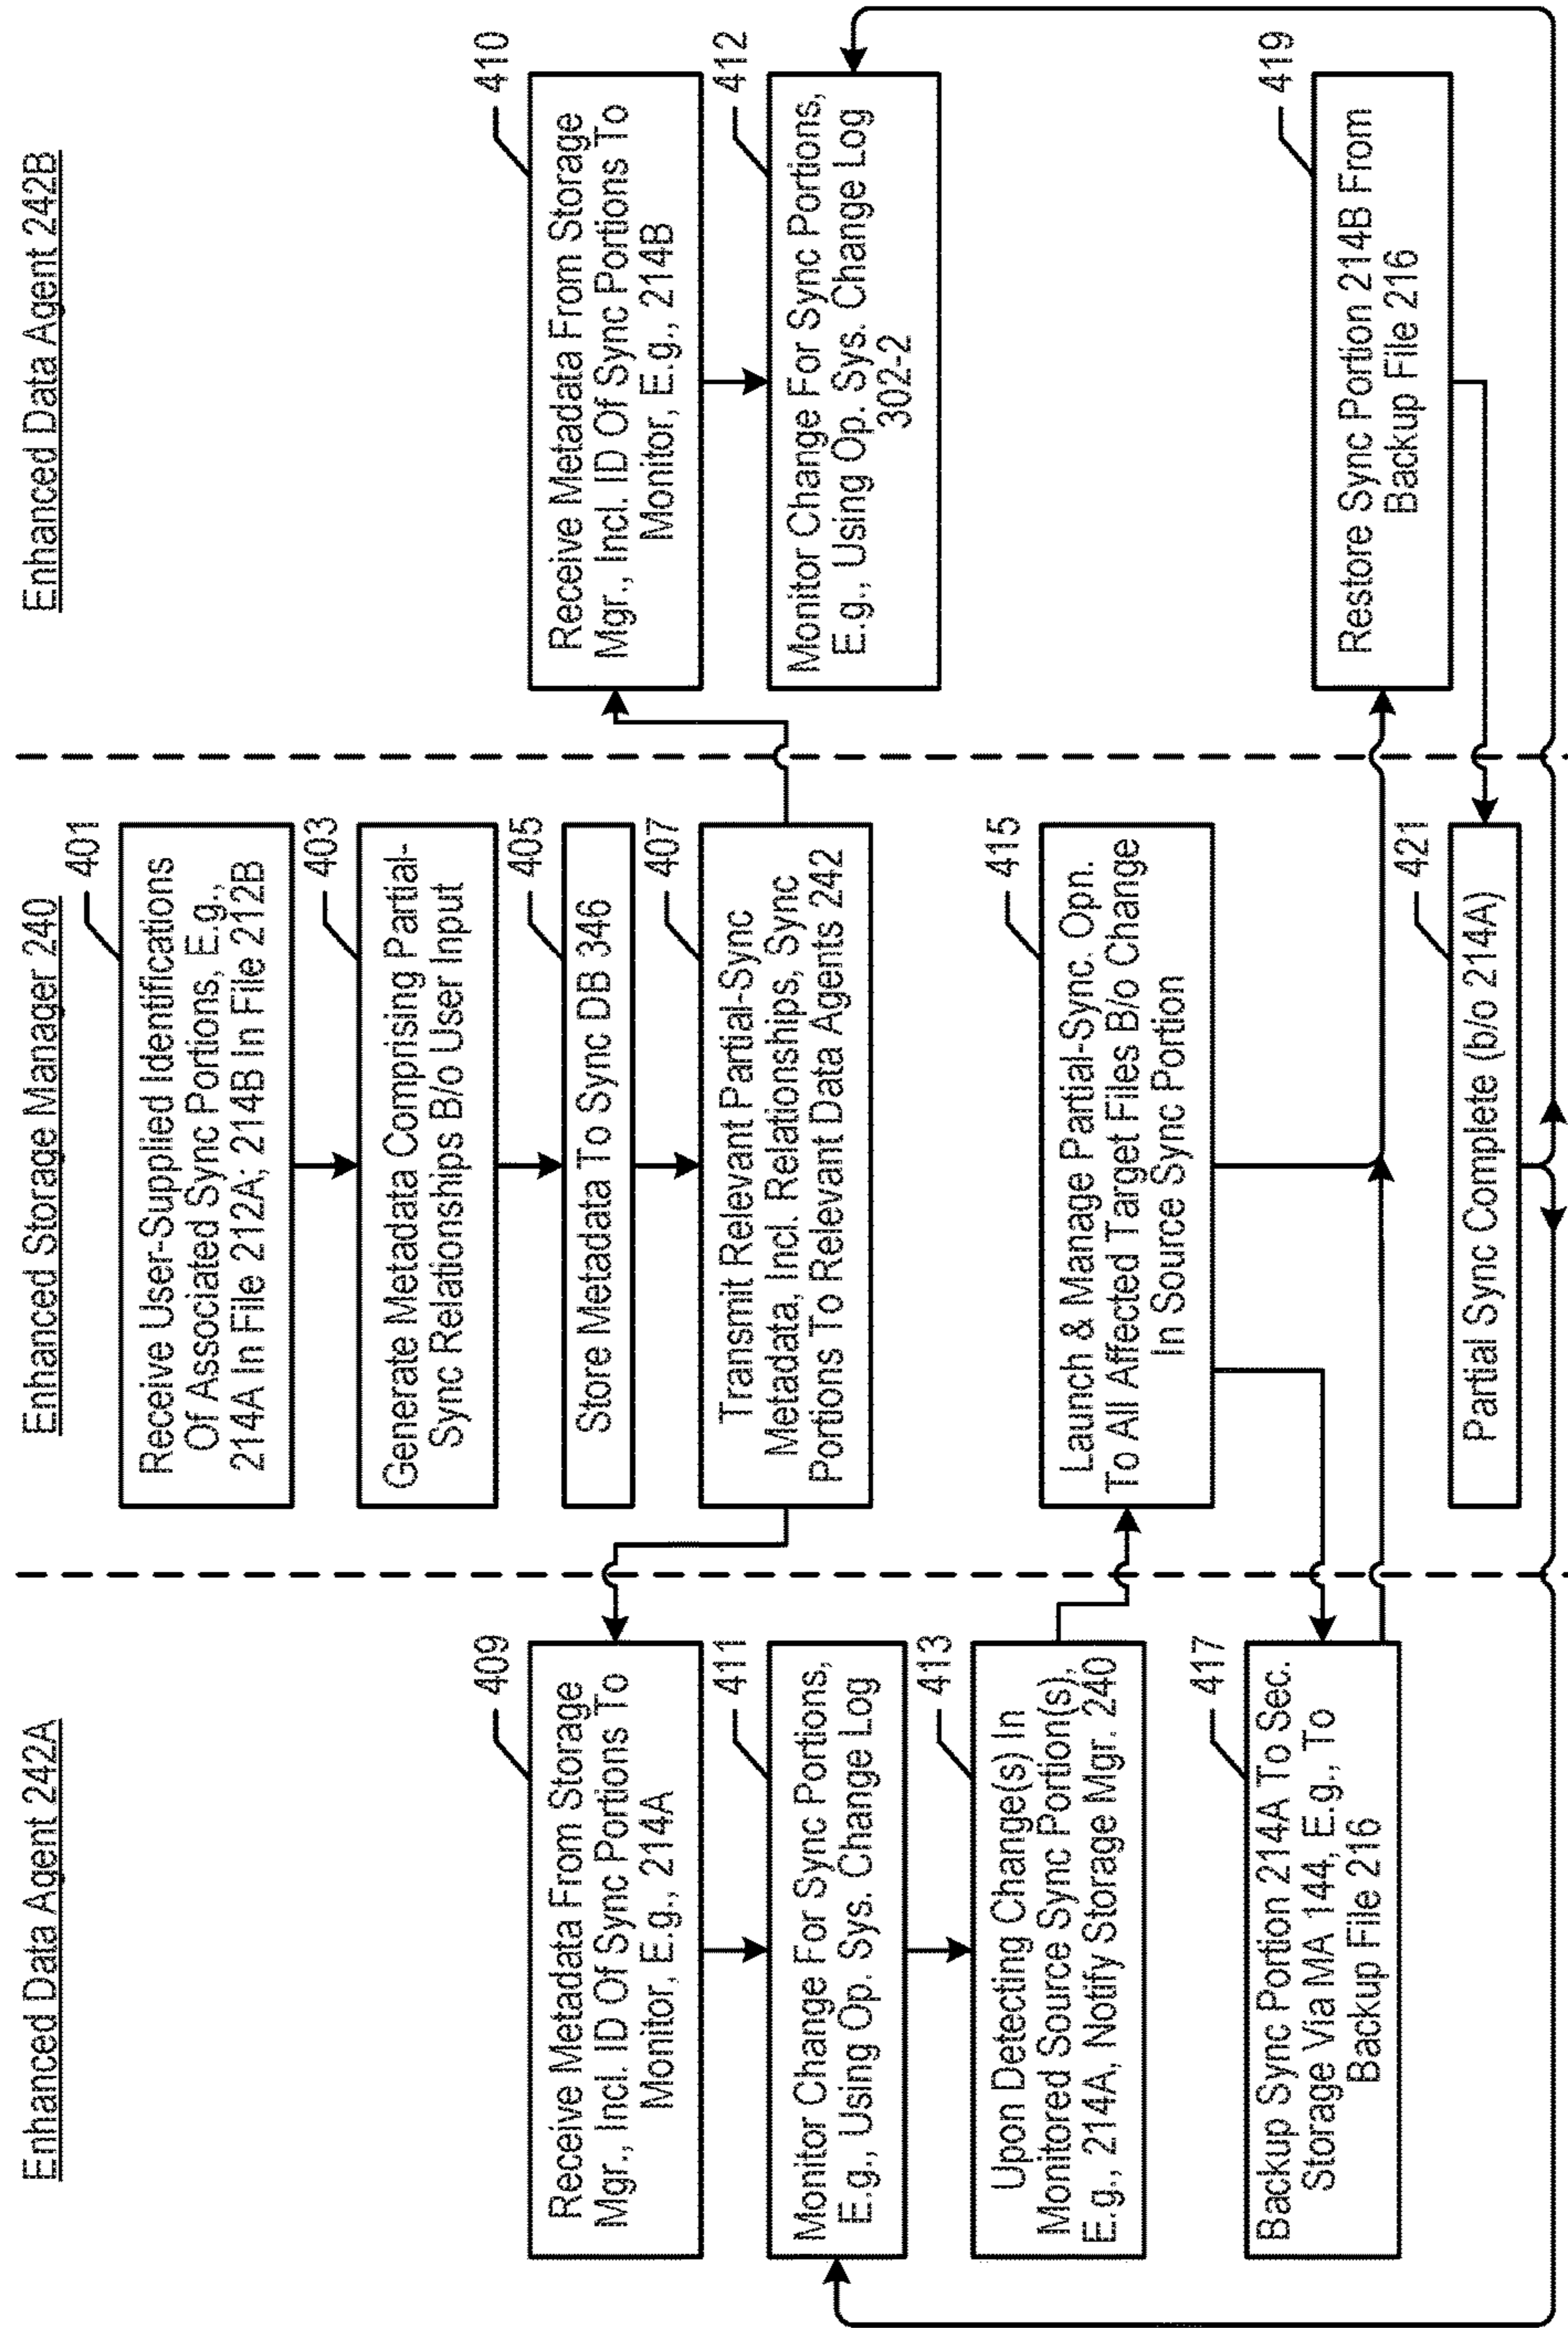


FIG. 4 Method 400 For Partially Synchronizing Primary Data Files Based On Synchronizing Portions (when 214A Changes)

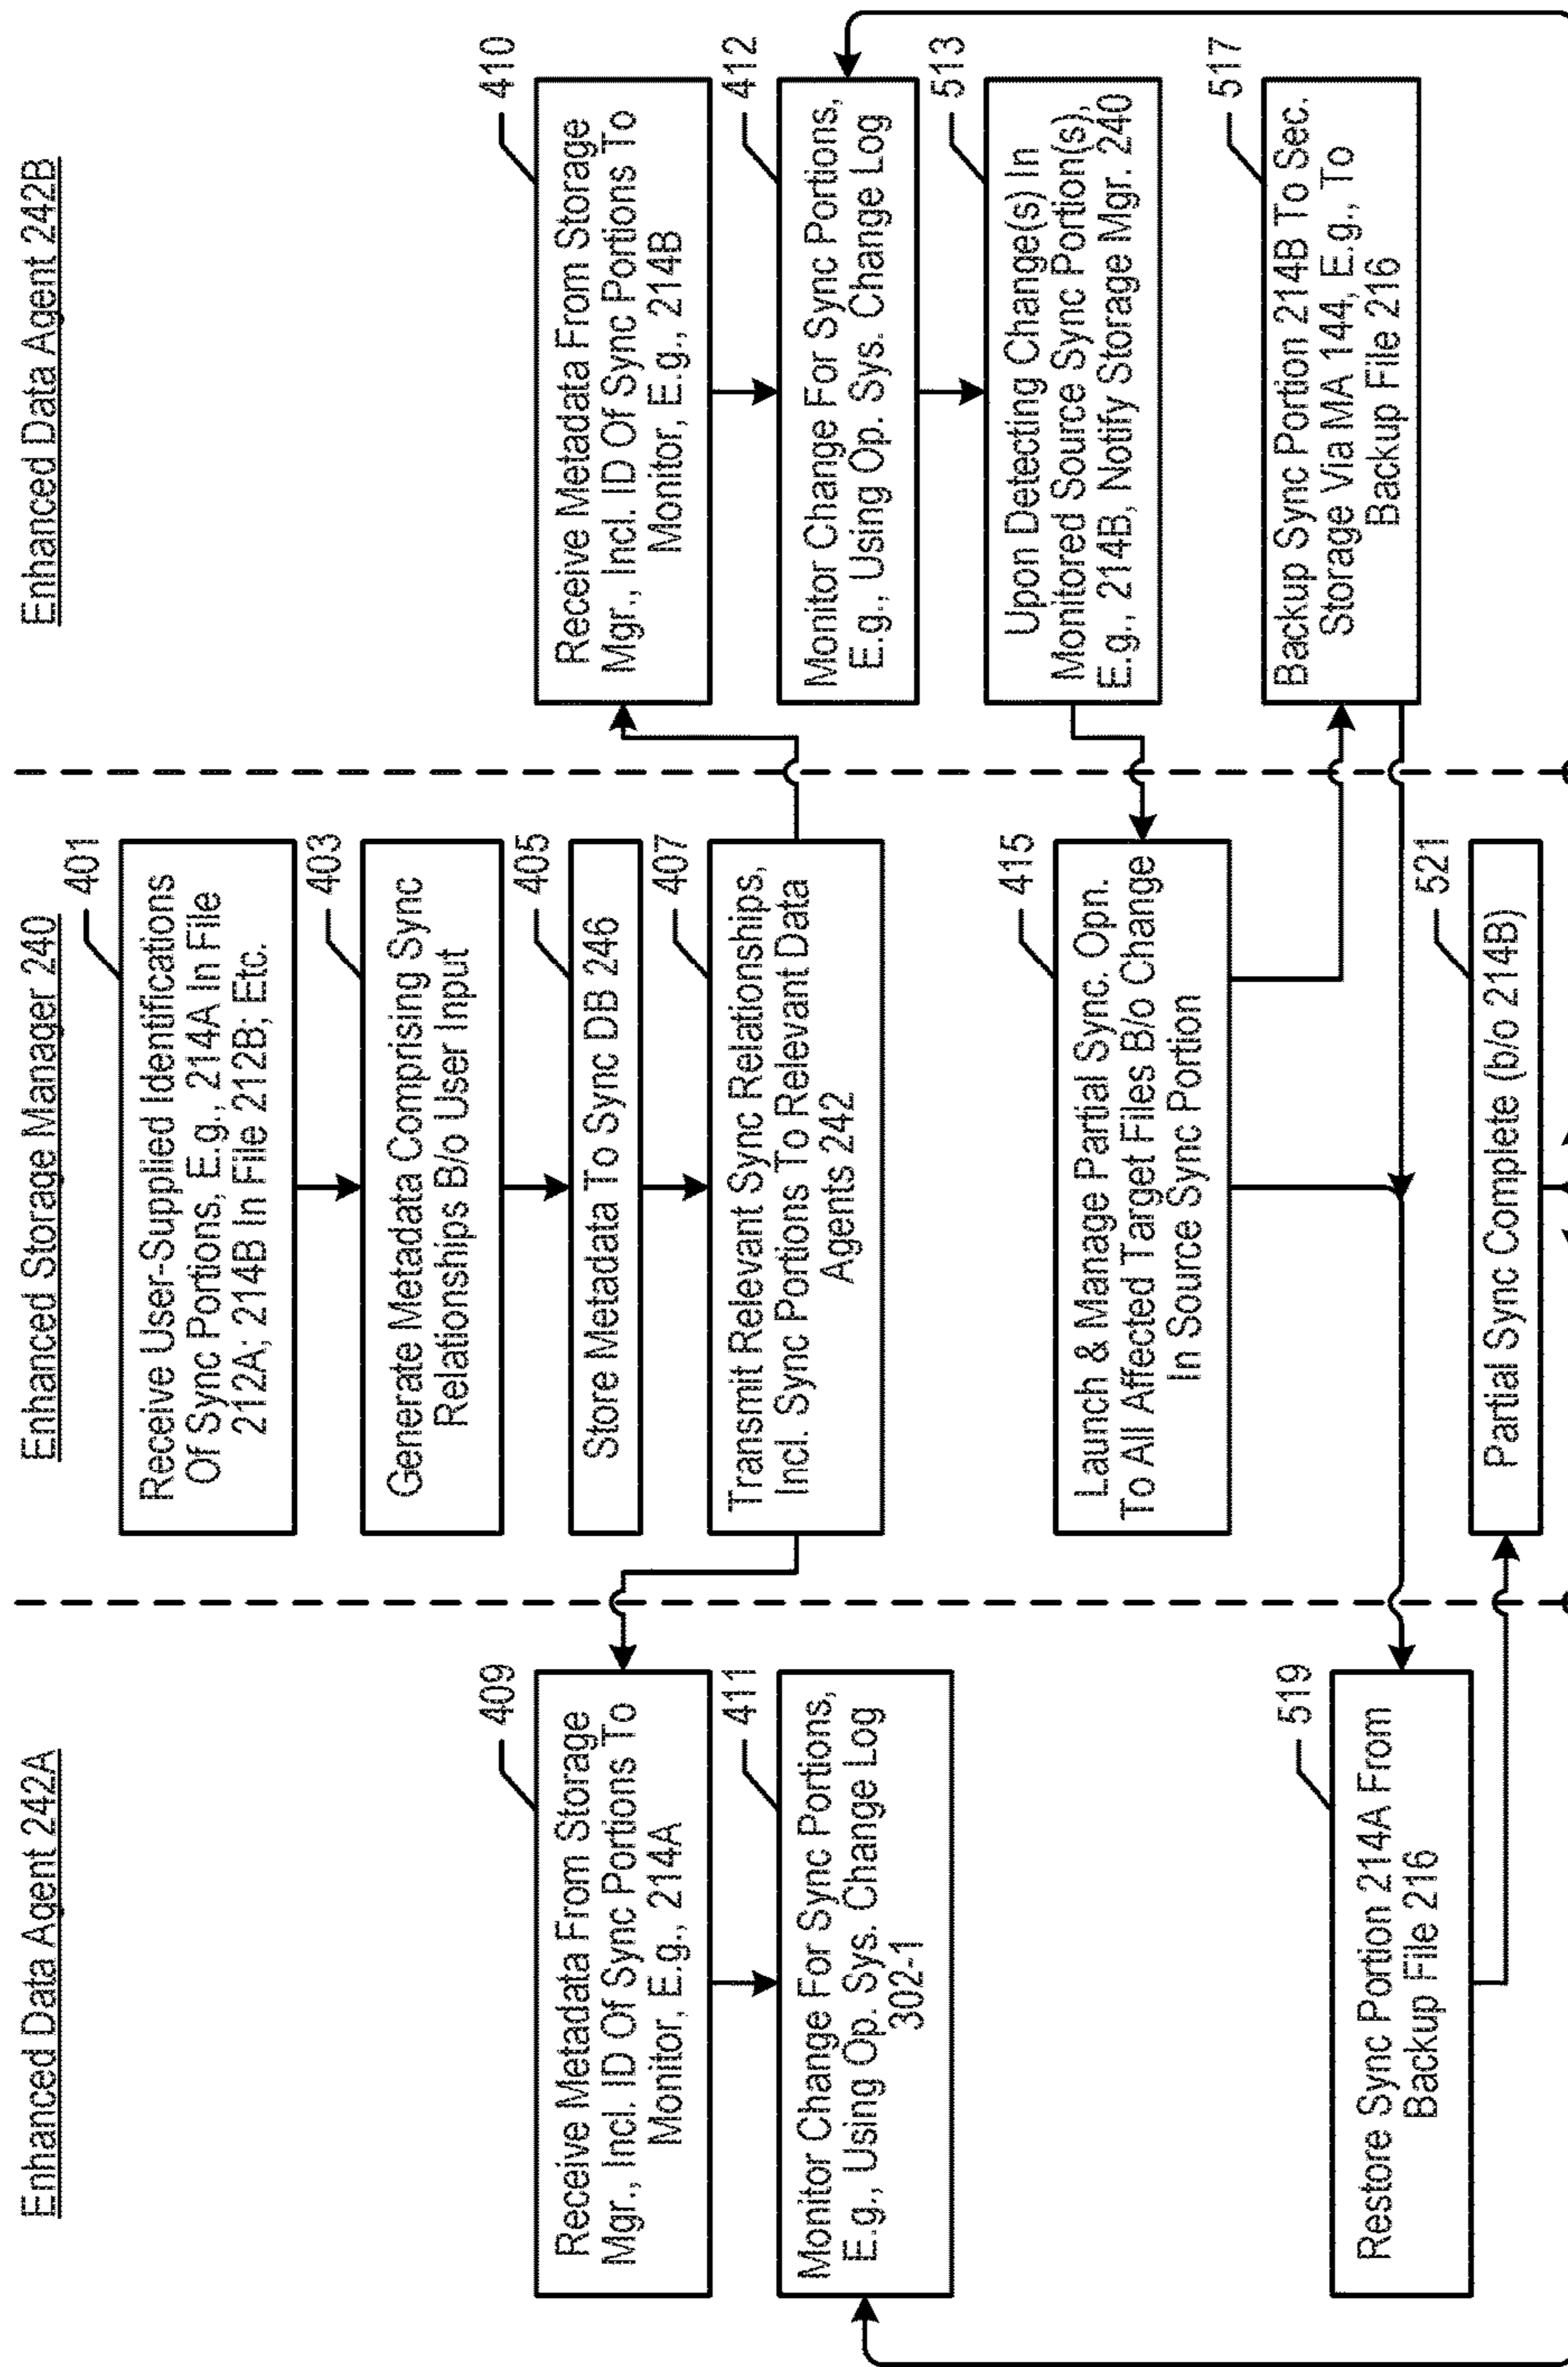


FIG. 5 Method 500 For Partially Synchronizing Primary Data Files Based On Synchronizing Portions (when 214B Changes)

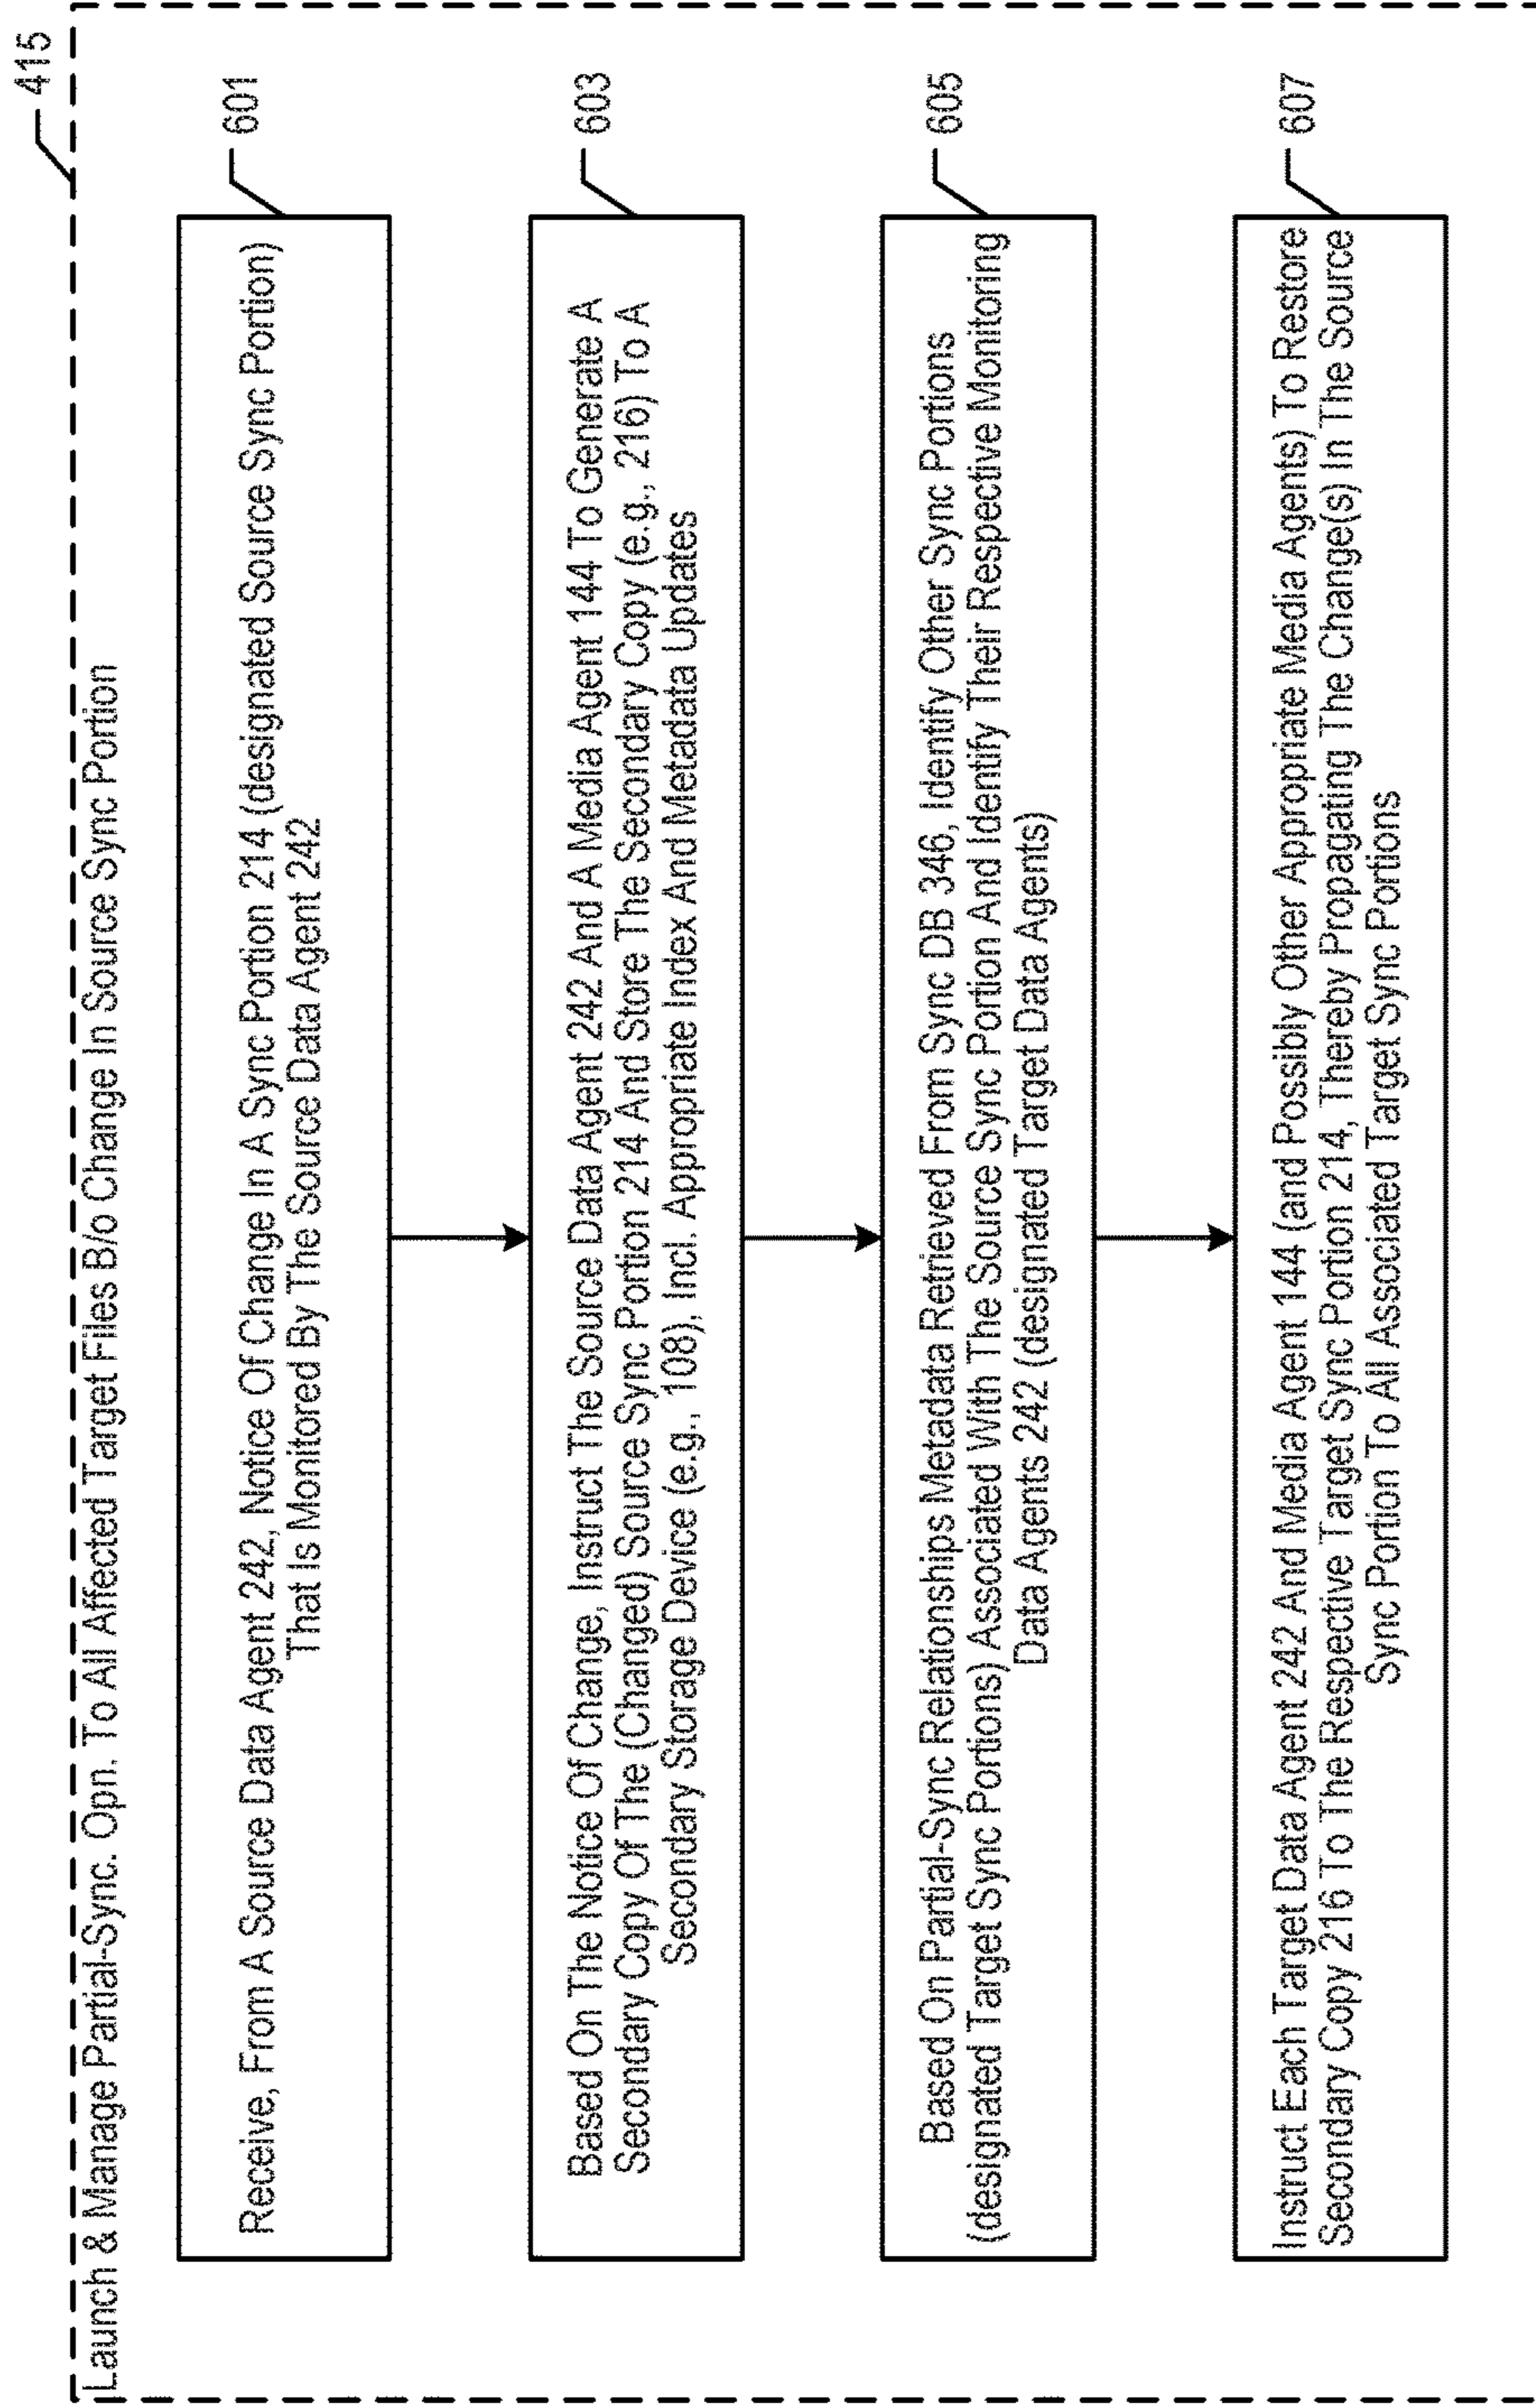


FIG. 6

1

SYNCHRONIZING SELECTED PORTIONS OF DATA IN A STORAGE MANAGEMENT SYSTEM

INCORPORATION BY REFERENCE TO ANY
PRIORITY APPLICATIONS

This application is a Continuation of U.S. patent application Ser. No. 15/412,645 entitled "Synchronizing Selected Portions of Data in a Storage Management System" and filed on Jan. 23, 2017, which is a Continuation of U.S. Pat. No. 9,588,849, entitled "Synchronizing Selected Portions of Data in a Storage Management System" and filed on Jan. 20, 2015. Any and all applications for which a foreign or domestic priority claim is identified in the Application Data Sheet of the present application are hereby incorporated by reference herein under 37 CFR 1.57.

BACKGROUND

Businesses worldwide recognize the commercial value of their data and seek reliable, cost-effective ways to protect the information stored on their computer networks while minimizing impact on productivity. A company might back up critical computing systems such as databases, virtual machines, file servers, web servers, and so on as part of a daily, weekly, or monthly maintenance schedule. Given the rapidly expanding volume of data under management, companies also continue to seek innovative techniques for managing data growth. In collaboration environments, a need arises to efficiently manage data growth and coordinate data content across groups of users.

SUMMARY

In collaboration environments, is desirable to find a way to coordinate data content, e.g., files, databases, etc., across a plurality of users and/or storage devices in such a way that network communications bandwidth and storage resources are kept relatively modest. Several users may wish to keep current with a rapidly changing source of data, e.g., a software development code base, a live database comprising transactions, a media project under development, etc. The traditional approach frequently copies the source data to a local workspace for the user's own use as needed. However, this approach is costly and may even be impractical when the source of data is large and numerous users wish to stay current. The traditional approach may cause network communication bottlenecks and may also use up too much storage and may be unsuitable for multi-party updates. The traditional approach also may affect performance in the production environment, by too-frequently tapping it for copies. Therefore, a more efficient and streamlined approach is needed. Even when changes are propagated incrementally, this may cover a much broader scope of changes than strictly necessary.

The present inventor devised methods and systems that leverage resources in a storage management system to partially synchronize primary data files based on synchronizing selected portions thereof without regard to changes that may be occurring in other non-synchronized portions. Accordingly, a number of primary data files may be partially synchronized by synchronizing designated portions thereof via auto-restore operations from backup data. This approach relies on storage management resources to: designate portions of source data that are to be kept synchronized across any number of targets; detect changes to the designated

2

portions; back up changes to secondary storage; and distribute the changes from secondary storage to the associated targets with minimal impact to the primary data environment. The approach may be mutual, so that changes in any one of an associated group of source data files may be likewise detected, backed up, and distributed to the other members of the group. By managing synchronization in specially-designated portions, i.e., using partial synchronization, instead of attempting to synchronize files in their entirety, the present approach uses substantially fewer communications and storage resources. Accordingly, the present approach may be more economically practicable and may enable partial-synchronization operations to occur more frequently and provide more current data to the respective users.

The illustrative embodiment employs an enhanced storage manager and enhanced data agents that work with media agents in the storage management system to perform partial-synchronization operations across any number of primary data files. The enhanced storage manager may receive designations from users and/or administrators that identify which portions of which primary data files they wish the storage management system to keep synchronized on an on-going basis. The storage manager comprises enhancements for storing metadata for designations of the selected portions (hereinafter "synchronization portions") and their mutual association, as well as enhancements for managing partial-synchronization operations throughout the storage management system. For example, the enhanced storage manager may receive notice that a synchronization portion of a primary data file has changed (e.g., a new transaction added to a particular section of a database, a change to a portion of source code, a change to a media file, etc.) and may then automatically launch a partial-synchronization operation to bring all the associated synchronization portions up-to-date with the detected change.

An enhanced data agent may comprise a monitor function that keeps track of the designated synchronization portion of a given data file that is under the purview of the data agent. When the enhanced data agent detects a change in the monitored synchronization portion, it may notify the storage manager. In response, the storage manager may instruct the data agent to back up the changed synchronization portion to secondary storage. Notably, only the synchronization portion is backed up, not the data file as a whole. The storage manager may then manage restore operations to restore the secondary copy to all the associated synchronization portions. No user interaction is required to pull in changes or synchronize the local copy of a data file with one or more associated other data files. The process may be managed by the enhanced storage manager, which instructs data agent and media agents to perform the requisite operations.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1A is a block diagram illustrating an exemplary information management system.

FIG. 1B is a detailed view of a primary storage device, a secondary storage device, and some examples of primary data and secondary copy data.

FIG. 1C is a block diagram of an exemplary information management system including a storage manager, one or more data agents, and one or more media agents.

FIG. 1D is a block diagram illustrating a scalable information management system.

FIG. 1E illustrates certain secondary copy operations according to an exemplary storage policy.

FIGS. 1F-1H are block diagrams illustrating suitable data structures that may be employed by the information management system.

FIG. 2 is a block diagram illustrating some salient portions of a storage management system **200** for partially synchronizing primary data files based on synchronizing portions thereof, according to an illustrative embodiment of the present invention.

FIG. 3 depicts some salient details of system **200** according to the illustrative embodiment.

FIG. 4 depicts some salient operations of a method **400** for partially synchronizing primary data files based on synchronizing portions thereof (e.g., when source synchronization portion **214A** changes) according to an illustrative embodiment of the present invention.

FIG. 5 depicts some salient operations of a method **500** for partially synchronizing primary data files based on synchronizing portions thereof (e.g., when source synchronization portion **214B** changes) according to the illustrative embodiment.

FIG. 6 depicts some salient sub-operations of block **415** in method **400** and method **500**.

DETAILED DESCRIPTION

Systems and methods are disclosed for partially synchronizing primary data files based on synchronizing portions thereof via auto-restore operations from backup data. Examples of such systems and methods are described in further detail herein, in reference to FIGS. 2-6. Components and functionality for partially synchronizing primary data files based on synchronizing portions thereof via auto-restore operations from backup data may be configured and/or incorporated into information management systems such as those described herein in FIGS. 1A-1H.

Information Management System Overview

With the increasing importance of protecting and leveraging data, organizations simply cannot afford to take the risk of losing critical data. Moreover, runaway data growth and other modern realities make protecting and managing data an increasingly difficult task. There is therefore a need for efficient, powerful, and user-friendly solutions for protecting and managing data.

Depending on the size of the organization, there are typically many data production sources which are under the purview of tens, hundreds, or even thousands of employees or other individuals. In the past, individual employees were sometimes responsible for managing and protecting their data. A patchwork of hardware and software point solutions has been applied in other cases. These solutions were often provided by different vendors and had limited or no interoperability.

Certain embodiments described herein provide systems and methods capable of addressing these and other shortcomings of prior approaches by implementing unified, organization-wide information management. FIG. 1A shows one such information management system **100**, which generally includes combinations of hardware and software configured to protect and manage data and metadata, which is generated and used by the various computing devices in information management system **100**. The organization that employs the information management system **100** may be a corporation or other business entity, non-profit organization, educational institution, household, governmental agency, or the like.

Generally, the systems and associated components described herein may be compatible with and/or provide some or all of the functionality of the systems and corre-

sponding components described in one or more of the following U.S. patents and patent application publications assigned to CommVault Systems, Inc., each of which is hereby incorporated in its entirety by reference herein:

U.S. Pat. No. 7,035,880, entitled "Modular Backup and Retrieval System Used in Conjunction With a Storage Area Network";

U.S. Pat. No. 7,107,298, entitled "System And Method For Archiving Objects In An Information Store";

U.S. Pat. No. 7,246,207, entitled "System and Method for Dynamically Performing Storage Operations in a Computer Network";

U.S. Pat. No. 7,315,923, entitled "System And Method For Combining Data Streams In Pipelined Storage Operations In A Storage Network";

U.S. Pat. No. 7,343,453, entitled "Hierarchical Systems and Methods for Providing a Unified View of Storage Information";

U.S. Pat. No. 7,395,282, entitled "Hierarchical Backup and Retrieval System";

U.S. Pat. No. 7,529,782, entitled "System and Methods for Performing a Snapshot and for Restoring Data";

U.S. Pat. No. 7,617,262, entitled "System and Methods for Monitoring Application Data in a Data Replication System";

U.S. Pat. No. 7,747,579, entitled "Metabase for Facilitating Data Classification";

U.S. Pat. No. 8,156,086, entitled "Systems And Methods For Stored Data Verification";

U.S. Pat. No. 8,170,995, entitled "Method and System for Offline Indexing of Content and Classifying Stored Data";

U.S. Pat. No. 8,229,954, entitled "Managing Copies Of Data";

U.S. Pat. No. 8,230,195, entitled "System And Method For Performing Auxiliary Storage Operations";

U.S. Pat. No. 8,285,681, entitled "Data Object Store and Server for a Cloud Storage Environment, Including Data Deduplication and Data Management Across Multiple Cloud Storage Sites";

U.S. Pat. No. 8,307,177, entitled "Systems And Methods For Management Of Virtualization Data";

U.S. Pat. No. 8,364,652, entitled "Content-Aligned, Block-Based Deduplication";

U.S. Pat. No. 8,578,120, entitled "Block-Level Single Instancing";

U.S. Pat. Pub. No. 2006/0224846, entitled "System and Method to Support Single Instance Storage Operations";

U.S. Pat. Pub. No. 2009/0319534, entitled "Application-Aware and Remote Single Instance Data Management";

U.S. Pat. Pub. No. 2012/0150818, entitled "Client-Side Repository in a Networked Deduplicated Storage System";

U.S. Pat. Pub. No. 2012/0150826, entitled "Distributed Deduplicated Storage System"; and

US Pat. Pub. No. 2014/0201142, entitled "Partial Sharing of Secondary Storage Files in a Data Storage System."

The information management system **100** can include a variety of different computing devices. For instance, as will be described in greater detail herein, the information management system **100** can include one or more client computing devices **102** and secondary storage computing devices **106**.

Computing devices can include, without limitation, one or more: workstations, personal computers, desktop comput-

ers, or other types of generally fixed computing systems such as mainframe computers and minicomputers. Other computing devices can include mobile or portable computing devices, such as one or more laptops, tablet computers, personal data assistants, mobile phones (such as smartphones), and other mobile or portable computing devices such as embedded computers, set top boxes, vehicle-mounted devices, wearable computers, etc. Computing devices can include servers, such as mail servers, file servers, database servers, and web servers.

In some cases, a computing device includes virtualized and/or cloud computing resources. For instance, one or more virtual machines may be provided to the organization by a third-party cloud service vendor. Or, in some embodiments, computing devices can include one or more virtual machine(s) running on a physical host computing device (or “host machine”) operated by the organization. As one example, the organization may use one virtual machine as a database server and another virtual machine as a mail server, both virtual machines operating on the same host machine.

A virtual machine includes an operating system and associated virtual resources, and is hosted simultaneously with another operating system on a physical host computer (or host machine). A hypervisor (typically software, and also known in the art as a virtual machine monitor or a virtual machine manager or “VMM”) sits between the virtual machine and the hardware of the physical host machine. One example of hypervisor as virtualization software is ESX Server, by VMware, Inc. of Palo Alto, Calif.; other examples include Microsoft Virtual Server and Microsoft Windows Server Hyper-V, both by Microsoft Corporation of Redmond, Wash., and Sun xVM by Oracle America Inc. of Santa Clara, Calif. In some embodiments, the hypervisor may be firmware or hardware or a combination of software and/or firmware and/or hardware.

The hypervisor provides to each virtual operating system virtual resources, such as a virtual processor, virtual memory, a virtual network device, and a virtual disk. Each virtual machine has one or more virtual disks. The hypervisor typically stores the data of virtual disks in files on the file system of the physical host machine, called virtual machine disk files (in the case of VMware virtual servers) or virtual hard disk image files (in the case of Microsoft virtual servers). For example, VMware’s ESX Server provides the Virtual Machine File System (VMFS) for the storage of virtual machine disk files. A virtual machine reads data from and writes data to its virtual disk much the same way that an actual physical machine reads data from and writes data to an actual disk.

Examples of techniques for implementing information management techniques in a cloud computing environment are described in U.S. Pat. No. 8,285,681, which is incorporated by reference herein. Examples of techniques for implementing information management techniques in a virtualized computing environment are described in U.S. Pat. No. 8,307,177, also incorporated by reference herein.

The information management system **100** can also include a variety of storage devices, including primary storage devices **104** and secondary storage devices **108**, for example. Storage devices can generally be of any suitable type including, without limitation, disk drives, hard-disk arrays, semiconductor memory (e.g., solid state storage devices), network attached storage (NAS) devices, tape libraries or other magnetic, non-tape storage devices, optical media storage devices, DNA/RNA-based memory technology, combinations of the same, and the like. In some embodiments, storage devices can form part of a distributed

file system. In some cases, storage devices are provided in a cloud (e.g., a private cloud or one operated by a third-party vendor). A storage device in some cases comprises a disk array or portion thereof.

The illustrated information management system **100** includes one or more client computing device **102** having at least one application **110** executing thereon, and one or more primary storage devices **104** storing primary data **112**. The client computing device(s) **102** and the primary storage devices **104** may generally be referred to in some cases as a primary storage subsystem **117**. A computing device in an information management system **100** that has a data agent **142** installed and operating on it is generally referred to as a client computing device **102** (or, in the context of a component of the information management system **100** simply as a “client”).

Depending on the context, the term “information management system” can refer to generally all of the illustrated hardware and software components. Or, in other instances, the term may refer to only a subset of the illustrated components.

For instance, in some cases, the information management system **100** generally refers to a combination of specialized components used to protect, move, manage, manipulate, analyze, and/or process data and metadata generated by the client computing devices **102**. However, the information management system **100** in some cases does not include the underlying components that generate and/or store the primary data **112**, such as the client computing devices **102** themselves, the applications **110** and operating system operating on the client computing devices **102**, and the primary storage devices **104**. As an example, “information management system” may sometimes refer to one or more of the following components and corresponding data structures: storage managers, data agents, and media agents. These components will be described in further detail below.

Client Computing Devices

There are typically a variety of sources in an organization that produce data to be protected and managed. As just one illustrative example, in a corporate environment such data sources can be employee workstations and company servers such as a mail server, a web server, a database server, a transaction server, or the like. In the information management system **100**, the data generation sources include the one or more client computing devices **102**.

The client computing devices **102** may include any of the types of computing devices described above, without limitation, and in some cases the client computing devices **102** are associated with one or more users and/or corresponding user accounts, of employees or other individuals.

The information management system **100** generally addresses and handles the data management and protection needs for the data generated by the client computing devices **102**. However, the use of this term does not imply that the client computing devices **102** cannot be “servers” in other respects. For instance, a particular client computing device **102** may act as a server with respect to other devices, such as other client computing devices **102**. As just a few examples, the client computing devices **102** can include mail servers, file servers, database servers, and web servers.

Each client computing device **102** may have one or more applications **110** (e.g., software applications) executing thereon which generate and manipulate the data that is to be protected from loss and managed. The applications **110** generally facilitate the operations of an organization (or multiple affiliated organizations), and can include, without limitation, mail server applications (e.g., Microsoft

Exchange Server), file server applications, mail client applications (e.g., Microsoft Exchange Client), database applications (e.g., SQL, Oracle, SAP, Lotus Notes Database), word processing applications (e.g., Microsoft Word), spreadsheet applications, financial applications, presentation applications, graphics and/or video applications, browser applications, mobile applications, entertainment applications, and so on.

The client computing devices **102** can have at least one operating system (e.g., Microsoft Windows, Mac OS X, iOS, IBM z/OS, Linux, other Unix-based operating systems, etc.) installed thereon, which may support or host one or more file systems and other applications **110**.

The client computing devices **102** and other components in information management system **100** can be connected to one another via one or more communication pathways **114**. For example, a first communication pathway **114** may connect (or communicatively couple) client computing device **102** and secondary storage computing device **106**; a second communication pathway **114** may connect storage manager **140** and client computing device **102**; and a third communication pathway **114** may connect storage manager **140** and secondary storage computing device **106**, etc. (see, e.g., FIG. 1A and FIG. 1C). The communication pathways **114** can include one or more networks or other connection types including one or more of the following, without limitation: the Internet, a wide area network (WAN), a local area network (LAN), a Storage Area Network (SAN), a Fibre Channel connection, a Small Computer System Interface (SCSI) connection, a virtual private network (VPN), a token ring or TCP/IP based network, an intranet network, a point-to-point link, a cellular network, a wireless data transmission system, a two-way cable system, an interactive kiosk network, a satellite network, a broadband network, a baseband network, a neural network, a mesh network, an ad hoc network, other appropriate wired, wireless, or partially wired/wireless computer or telecommunications networks, combinations of the same or the like. The communication pathways **114** in some cases may also include application programming interfaces (APIs) including, e.g., cloud service provider APIs, virtual machine management APIs, and hosted service provider APIs. The underlying infrastructure of communication paths **114** may be wired and/or wireless, analog and/or digital, or any combination thereof; and the facilities used may be private, public, third-party provided, or any combination thereof, without limitation.

Primary Data and Exemplary Primary Storage Devices

Primary data **112** according to some embodiments is production data or other “live” data generated by the operating system and/or applications **110** operating on a client computing device **102**. The primary data **112** is generally stored on the primary storage device(s) **104** and is organized via a file system supported by the client computing device **102**. For instance, the client computing device(s) **102** and corresponding applications **110** may create, access, modify, write, delete, and otherwise use primary data **112**. In some cases, some or all of the primary data **112** can be stored in cloud storage resources (e.g., primary storage device **104** may be a cloud-based resource).

Primary data **112** is generally in the native format of the source application **110**. According to certain aspects, primary data **112** is an initial or first (e.g., created before any other copies or before at least one other copy) stored copy of data generated by the source application **110**. Primary data **112** in some cases is created substantially directly from data generated by the corresponding source applications **110**.

The primary storage devices **104** storing the primary data **112** may be relatively fast and/or expensive technology (e.g., a disk drive, a hard-disk array, solid state memory, etc.). In addition, primary data **112** may be highly changeable and/or may be intended for relatively short term retention (e.g., hours, days, or weeks).

According to some embodiments, the client computing device **102** can access primary data **112** from the primary storage device **104** by making conventional file system calls via the operating system. Primary data **112** may include structured data (e.g., database files), unstructured data (e.g., documents), and/or semi-structured data. Some specific examples are described below with respect to FIG. 1B.

It can be useful in performing certain tasks to organize the primary data **112** into units of different granularities. In general, primary data **112** can include files, directories, file system volumes, data blocks, extents, or any other hierarchies or organizations of data objects. As used herein, a “data object” can refer to both (1) any file that is currently addressable by a file system or that was previously addressable by the file system (e.g., an archive file) and (2) a subset of such a file (e.g., a data block).

As will be described in further detail, it can also be useful in performing certain functions of the information management system **100** to access and modify metadata within the primary data **112**. Metadata generally includes information about data objects or characteristics associated with the data objects. For simplicity herein, it is to be understood that, unless expressly stated otherwise, any reference to primary data **112** generally also includes its associated metadata, but references to the metadata do not include the primary data.

Metadata can include, without limitation, one or more of the following: the data owner (e.g., the client or user that generates the data), the last modified time (e.g., the time of the most recent modification of the data object), a data object name (e.g., a file name), a data object size (e.g., a number of bytes of data), information about the content (e.g., an indication as to the existence of a particular search term), user-supplied tags, to/from information for email (e.g., an email sender, recipient, etc.), creation date, file type (e.g., format or application type), last accessed time, application type (e.g., type of application that generated the data object), location/network (e.g., a current, past or future location of the data object and network pathways to/from the data object), geographic location (e.g., GPS coordinates), frequency of change (e.g., a period in which the data object is modified), business unit (e.g., a group or department that generates, manages or is otherwise associated with the data object), aging information (e.g., a schedule, such as a time period, in which the data object is migrated to secondary or long term storage), boot sectors, partition layouts, file location within a file folder directory structure, user permissions, owners, groups, access control lists [ACLs]), system metadata (e.g., registry information), combinations of the same or other similar information related to the data object.

In addition to metadata generated by or related to file systems and operating systems, some of the applications **110** and/or other components of the information management system **100** maintain indices of metadata for data objects, e.g., metadata associated with individual email messages. Thus, each data object may be associated with corresponding metadata. The use of metadata to perform classification and other functions is described in greater detail below.

Each of the client computing devices **102** are generally associated with and/or in communication with one or more of the primary storage devices **104** storing corresponding primary data **112**. A client computing device **102** may be

considered to be “associated with” or “in communication with” a primary storage device **104** if it is capable of one or more of: routing and/or storing data (e.g., primary data **112**) to the particular primary storage device **104**, coordinating the routing and/or storing of data to the particular primary storage device **104**, retrieving data from the particular primary storage device **104**, coordinating the retrieval of data from the particular primary storage device **104**, and modifying and/or deleting data retrieved from the particular primary storage device **104**.

The primary storage devices **104** can include any of the different types of storage devices described above, or some other kind of suitable storage device. The primary storage devices **104** may have relatively fast I/O times and/or are relatively expensive in comparison to the secondary storage devices **108**. For example, the information management system **100** may generally regularly access data and metadata stored on primary storage devices **104**, whereas data and metadata stored on the secondary storage devices **108** is accessed relatively less frequently.

Primary storage device **104** may be dedicated or shared. In some cases, each primary storage device **104** is dedicated to an associated client computing device **102**. For instance, a primary storage device **104** in one embodiment is a local disk drive of a corresponding client computing device **102**. In other cases, one or more primary storage devices **104** can be shared by multiple client computing devices **102**, e.g., via a network such as in a cloud storage implementation. As one example, a primary storage device **104** can be a disk array shared by a group of client computing devices **102**, such as one of the following types of disk arrays: EMC Clariion, EMC Symmetrix, EMC Celerra, Dell EqualLogic, IBM XIV, NetApp FAS, HP EVA, and HP 3PAR.

The information management system **100** may also include hosted services (not shown), which may be hosted in some cases by an entity other than the organization that employs the other components of the information management system **100**. For instance, the hosted services may be provided by various online service providers to the organization. Such service providers can provide services including social networking services, hosted email services, or hosted productivity applications or other hosted applications). Hosted services may include software-as-a-service (SaaS), platform-as-a-service (PaaS), application service providers (ASPs), cloud services, or other mechanisms for delivering functionality via a network. As it provides services to users, each hosted service may generate additional data and metadata under management of the information management system **100**, e.g., as primary data **112**. In some cases, the hosted services may be accessed using one of the applications **110**. As an example, a hosted mail service may be accessed via browser running on a client computing device **102**. The hosted services may be implemented in a variety of computing environments. In some cases, they are implemented in an environment having a similar arrangement to the information management system **100**, where various physical and logical components are distributed over a network.

Secondary Copies and Exemplary Secondary Storage Devices

The primary data **112** stored on the primary storage devices **104** may be compromised in some cases, such as when an employee deliberately or accidentally deletes or overwrites primary data **112** during their normal course of work. Or the primary storage devices **104** can be damaged, lost, or otherwise corrupted. For recovery and/or regulatory compliance purposes, it is therefore useful to generate copies

of the primary data **112**. Accordingly, the information management system **100** includes one or more secondary storage computing devices **106** and one or more secondary storage devices **108** configured to create and store one or more secondary copies **116** of the primary data **112** and associated metadata. The secondary storage computing devices **106** and the secondary storage devices **108** may sometimes be referred to as a secondary storage subsystem **118**.

Creation of secondary copies **116** can help in search and analysis efforts and meet other information management goals, such as: restoring data and/or metadata if an original version (e.g., of primary data **112**) is lost (e.g., by deletion, corruption, or disaster); allowing point-in-time recovery; complying with regulatory data retention and electronic discovery (e-discovery) requirements; reducing utilized storage capacity; facilitating organization and search of data; improving user access to data files across multiple computing devices and/or hosted services; and implementing data retention policies.

The client computing devices **102** access or receive primary data **112** and communicate the data, e.g., over one or more communication pathways **114**, for storage in the secondary storage device(s) **108**.

A secondary copy **116** can comprise a separate stored copy of application data that is derived from one or more earlier-created, stored copies (e.g., derived from primary data **112** or another secondary copy **116**). Secondary copies **116** can include point-in-time data, and may be intended for relatively long-term retention (e.g., weeks, months or years), before some or all of the data is moved to other storage or is discarded.

In some cases, a secondary copy **116** is a copy of application data created and stored subsequent to at least one other stored instance (e.g., subsequent to corresponding primary data **112** or to another secondary copy **116**), in a different storage device than at least one previous stored copy, and/or remotely from at least one previous stored copy. In some other cases, secondary copies can be stored in the same storage device as primary data **112** and/or other previously stored copies. For example, in one embodiment a disk array capable of performing hardware snapshots stores primary data **112** and creates and stores hardware snapshots of the primary data **112** as secondary copies **116**. Secondary copies **116** may be stored in relatively slow and/or low cost storage (e.g., magnetic tape). A secondary copy **116** may be stored in a backup or archive format, or in some other format different than the native source application format or other primary data format.

In some cases, secondary copies **116** are indexed so users can browse and restore at another point in time. After creation of a secondary copy **116** representative of certain primary data **112**, a pointer or other location indicia (e.g., a stub) may be placed in primary data **112**, or be otherwise associated with primary data **112** to indicate the current location on the secondary storage device(s) **108** of secondary copy **116**.

Since an instance of a data object or metadata in primary data **112** may change over time as it is modified by an application **110** (or hosted service or the operating system), the information management system **100** may create and manage multiple secondary copies **116** of a particular data object or metadata, each representing the state of the data object in primary data **112** at a particular point in time. Moreover, since an instance of a data object in primary data **112** may eventually be deleted from the primary storage device **104** and the file system, the information management system **100** may continue to manage point-in-time repre-

11

sentations of that data object, even though the instance in primary data **112** no longer exists.

For virtualized computing devices the operating system and other applications **110** of the client computing device(s) **102** may execute within or under the management of virtualization software (e.g., a VMM), and the primary storage device(s) **104** may comprise a virtual disk created on a physical storage device. The information management system **100** may create secondary copies **116** of the files or other data objects in a virtual disk file and/or secondary copies **116** of the entire virtual disk file itself (e.g., of an entire .vmdk file).

Secondary copies **116** may be distinguished from corresponding primary data **112** in a variety of ways, some of which will now be described. First, as discussed, secondary copies **116** can be stored in a different format (e.g., backup, archive, or other non-native format) than primary data **112**. For this or other reasons, secondary copies **116** may not be directly useable by the applications **110** of the client computing device **102**, e.g., via standard system calls or otherwise without modification, processing, or other intervention by the information management system **100**.

Secondary copies **116** are also in some embodiments stored on a secondary storage device **108** that is inaccessible to the applications **110** running on the client computing devices **102** (and/or hosted services). Some secondary copies **116** may be “offline copies,” in that they are not readily available (e.g., not mounted to tape or disk). Offline copies can include copies of data that the information management system **100** can access without human intervention (e.g., tapes within an automated tape library, but not yet mounted in a drive), and copies that the information management system **100** can access only with at least some human intervention (e.g., tapes located at an offsite storage site).
The Use of Intermediate Devices for Creating Secondary Copies

Creating secondary copies can be a challenging task. For instance, there can be hundreds or thousands of client computing devices **102** continually generating large volumes of primary data **112** to be protected. Also, there can be significant overhead involved in the creation of secondary copies **116**. Moreover, secondary storage devices **108** may be special purpose components, and interacting with them can require specialized intelligence.

In some cases, the client computing devices **102** interact directly with the secondary storage device **108** to create the secondary copies **116**. However, in view of the factors described above, this approach can negatively impact the ability of the client computing devices **102** to serve the applications **110** and produce primary data **112**. Further, the client computing devices **102** may not be optimized for interaction with the secondary storage devices **108**.

Thus, in some embodiments, the information management system **100** includes one or more software and/or hardware components which generally act as intermediaries between the client computing devices **102** and the secondary storage devices **108**. In addition to off-loading certain responsibilities from the client computing devices **102**, these intermediate components can provide other benefits. For instance, as discussed further below with respect to FIG. 1D, distributing some of the work involved in creating secondary copies **116** can enhance scalability.

The intermediate components can include one or more secondary storage computing devices **106** as shown in FIG. 1A and/or one or more media agents, which can be software modules operating on corresponding secondary storage

12

computing devices **106** (or other appropriate computing devices). Media agents are discussed below (e.g., with respect to FIGS. 1C-1E).

The secondary storage computing device(s) **106** can comprise any of the computing devices described above, without limitation. In some cases, the secondary storage computing device(s) **106** include specialized hardware and/or software componentry for interacting with the secondary storage devices **108**.

To create a secondary copy **116** involving the copying of data from the primary storage subsystem **117** to the secondary storage subsystem **118**, the client computing device **102** in some embodiments communicates the primary data **112** to be copied (or a processed version thereof) to the designated secondary storage computing device **106**, via the communication pathway **114**. The secondary storage computing device **106** in turn conveys the received data (or a processed version thereof) to the secondary storage device **108**. In some such configurations, the communication pathway **114** between the client computing device **102** and the secondary storage computing device **106** comprises a portion of a LAN, WAN or SAN. In other cases, at least some client computing devices **102** communicate directly with the secondary storage devices **108** (e.g., via Fibre Channel or SCSI connections). In some other cases, one or more secondary copies **116** are created from existing secondary copies, such as in the case of an auxiliary copy operation, described in greater detail below.

Exemplary Primary Data and an Exemplary Secondary Copy

FIG. 1B is a detailed view showing some specific examples of primary data stored on the primary storage device(s) **104** and secondary copy data stored on the secondary storage device(s) **108**, with other components in the system removed for the purposes of illustration. Stored on the primary storage device(s) **104** are primary data objects including word processing documents **119A-B**, spreadsheets **120**, presentation documents **122**, video files **124**, image files **126**, email mailboxes **128** (and corresponding email messages **129A-C**), html/xml or other types of markup language files **130**, databases **132** and corresponding tables or other data structures **133A-133C**.

Some or all primary data objects are associated with corresponding metadata (e.g., “Meta1-11”), which may include file system metadata and/or application specific metadata. Stored on the secondary storage device(s) **108** are secondary copy data objects **134A-C** which may include copies of or otherwise represent corresponding primary data objects and metadata.

As shown, the secondary copy data objects **134A-C** can individually represent more than one primary data object. For example, secondary copy data object **134A** represents three separate primary data objects **133C**, **122**, and **129C** (represented as **133C'**, **122'**, and **129C'**, respectively, and accompanied by the corresponding metadata Meta**11**, Meta**3**, and Meta**8**, respectively). Moreover, as indicated by the prime mark ('), a secondary copy object may store a representation of a primary data object and/or metadata differently than the original format, e.g., in a compressed, encrypted, deduplicated, or other modified format. Likewise, secondary data object **134B** represents primary data objects **120**, **133B**, and **119A** as **120'**, **133B'**, and **119A'**, respectively and accompanied by corresponding metadata Meta**2**, Meta**10**, and Meta**1**, respectively. Also, secondary data object **134C** represents primary data objects **133A**,

1196, and 129A as 133A', 1196', and 129A', respectively, accompanied by corresponding metadata Meta9, Meta5, and Meta6, respectively.

Exemplary Information Management System Architecture

The information management system 100 can incorporate a variety of different hardware and software components, which can in turn be organized with respect to one another in many different configurations, depending on the embodiment. There are critical design choices involved in specifying the functional responsibilities of the components and the role of each component in the information management system 100. For instance, as will be discussed, such design choices can impact performance as well as the adaptability of the information management system 100 to data growth or other changing circumstances.

FIG. 1C shows an information management system 100 designed according to these considerations and which includes: storage manager 140, a centralized storage and/or information manager that is configured to perform certain control functions, one or more data agents 142 executing on the client computing device(s) 102 configured to process primary data 112, and one or more media agents 144 executing on the one or more secondary storage computing devices 106 for performing tasks involving the secondary storage devices 108. While distributing functionality amongst multiple computing devices can have certain advantages, in other contexts it can be beneficial to consolidate functionality on the same computing device. As such, in various other embodiments, one or more of the components shown in FIG. 1C as being implemented on separate computing devices are implemented on the same computing device. In one configuration, a storage manager 140, one or more data agents 142, and one or more media agents 144 are all implemented on the same computing device. In another embodiment, one or more data agents 142 and one or more media agents 144 are implemented on the same computing device, while the storage manager 140 is implemented on a separate computing device, etc. without limitation.

Storage Manager

As noted, the number of components in the information management system 100 and the amount of data under management can be quite large. Managing the components and data is therefore a significant task, and a task that can grow in an often unpredictable fashion as the quantity of components and data scale to meet the needs of the organization. For these and other reasons, according to certain embodiments, responsibility for controlling the information management system 100, or at least a significant portion of that responsibility, is allocated to the storage manager 140. By distributing control functionality in this manner, the storage manager 140 can be adapted independently according to changing circumstances. Moreover, a computing device for hosting the storage manager 140 can be selected to best suit the functions of the storage manager 140. These and other advantages are described in further detail below with respect to FIG. 1D.

The storage manager 140 may be a software module or other application, which, in some embodiments operates in conjunction with one or more associated data structures, e.g., a dedicated database (e.g., management database 146). In some embodiments, storage manager 140 is a computing device comprising circuitry for executing computer instructions and performs the functions described herein. The storage manager generally initiates, performs, coordinates and/or controls storage and other information management operations performed by the information management system 100, e.g., to protect and control the primary data 112 and

secondary copies 116 of data and metadata. In general, storage manager 100 may be said to manage information management system 100, which includes managing the constituent components, e.g., data agents and media agents, etc.

As shown by the dashed arrowed lines 114 in FIG. 1C, the storage manager 140 may communicate with and/or control some or all elements of the information management system 100, such as the data agents 142 and media agents 144. Thus, in certain embodiments, control information originates from the storage manager 140 and status reporting is transmitted to storage manager 140 by the various managed components, whereas payload data and payload metadata is generally communicated between the data agents 142 and the media agents 144 (or otherwise between the client computing device(s) 102 and the secondary storage computing device(s) 106), e.g., at the direction of and under the management of the storage manager 140. Control information can generally include parameters and instructions for carrying out information management operations, such as, without limitation, instructions to perform a task associated with an operation, timing information specifying when to initiate a task associated with an operation, data path information specifying what components to communicate with or access in carrying out an operation, and the like. Payload data, on the other hand, can include the actual data involved in the storage operation, such as content data written to a secondary storage device 108 in a secondary copy operation. Payload metadata can include any of the types of metadata described herein, and may be written to a storage device along with the payload content data (e.g., in the form of a header).

In other embodiments, some information management operations are controlled by other components in the information management system 100 (e.g., the media agent(s) 144 or data agent(s) 142), instead of or in combination with the storage manager 140.

According to certain embodiments, the storage manager 140 provides one or more of the following functions:

- initiating execution of secondary copy operations;
- managing secondary storage devices 108 and inventory/capacity of the same;
- reporting, searching, and/or classification of data in the information management system 100;
- allocating secondary storage devices 108 for secondary storage operations;
- monitoring completion of and providing status reporting related to secondary storage operations;
- tracking age information relating to secondary copies 116, secondary storage devices 108, and comparing the age information against retention guidelines;
- tracking movement of data within the information management system 100;
- tracking logical associations between components in the information management system 100;
- protecting metadata associated with the information management system 100; and
- implementing operations management functionality.

The storage manager 140 may maintain a database 146 (or "storage manager database 146" or "management database 146") of management-related data and information management policies 148. The database 146 may include a management index 150 (or "index 150") or other data structure that stores logical associations between components of the system, user preferences and/or profiles (e.g., preferences regarding encryption, compression, or deduplication of primary or secondary copy data, preferences regarding the

scheduling, type, or other aspects of primary or secondary copy or other operations, mappings of particular information management users or user accounts to certain computing devices or other components, etc.), management tasks, media containerization, or other useful data. For example, the storage manager **140** may use the index **150** to track logical associations between media agents **144** and secondary storage devices **108** and/or movement of data from primary storage devices **104** to secondary storage devices **108**. For instance, the index **150** may store data associating a client computing device **102** with a particular media agent **144** and/or secondary storage device **108**, as specified in an information management policy **148** (e.g., a storage policy, which is defined in more detail below).

Administrators and other people may be able to configure and initiate certain information management operations on an individual basis. But while this may be acceptable for some recovery operations or other relatively less frequent tasks, it is often not workable for implementing on-going organization-wide data protection and management. Thus, the information management system **100** may utilize information management policies **148** for specifying and executing information management operations (e.g., on an automated basis). Generally, an information management policy **148** can include a data structure or other information source that specifies a set of parameters (e.g., criteria and rules) associated with storage or other information management operations.

The storage manager database **146** may maintain the information management policies **148** and associated data, although the information management policies **148** can be stored in any appropriate location. For instance, an information management policy **148** such as a storage policy may be stored as metadata in a media agent database **152** or in a secondary storage device **108** (e.g., as an archive copy) for use in restore operations or other information management operations, depending on the embodiment. Information management policies **148** are described further below.

According to certain embodiments, the storage manager database **146** comprises a relational database (e.g., an SQL database) for tracking metadata, such as metadata associated with secondary copy operations (e.g., what client computing devices **102** and corresponding data were protected). This and other metadata may additionally be stored in other locations, such as at the secondary storage computing devices **106** or on the secondary storage devices **108**, allowing data recovery without the use of the storage manager **140** in some cases.

As shown, the storage manager **140** may include a jobs agent **156**, a user interface **158**, and a management agent **154**, all of which may be implemented as interconnected software modules or application programs.

The jobs agent **156** in some embodiments initiates, controls, and/or monitors the status of some or all storage or other information management operations previously performed, currently being performed, or scheduled to be performed by the information management system **100**. For instance, the jobs agent **156** may access information management policies **148** to determine when and how to initiate and control secondary copy and other information management operations, as will be discussed further.

The user interface **158** may include information processing and display software, such as a graphical user interface (“GUI”), an application program interface (“API”), or other interactive interface(s) through which users and system processes can retrieve information about the status of information management operations (e.g., storage operations) or

issue instructions to the information management system **100** and its constituent components. Via the user interface **158**, users may optionally issue instructions to the components in the information management system **100** regarding performance of storage and recovery operations. For example, a user may modify a schedule concerning the number of pending secondary copy operations. As another example, a user may employ the GUI to view the status of pending storage operations or to monitor the status of certain components in the information management system **100** (e.g., the amount of capacity left in a storage device).

An “information management cell” (or “storage operation cell” or “cell”) may generally include a logical and/or physical grouping of a combination of hardware and software components associated with performing information management operations on electronic data, typically one storage manager **140** and at least one client computing device **102** (comprising data agent(s) **142**) and at least one media agent **144**. For instance, the components shown in FIG. **1C** may together form an information management cell. Multiple cells may be organized hierarchically. With this configuration, cells may inherit properties from hierarchically superior cells or be controlled by other cells in the hierarchy (automatically or otherwise). Alternatively, in some embodiments, cells may inherit or otherwise be associated with information management policies, preferences, information management metrics, or other properties or characteristics according to their relative position in a hierarchy of cells. Cells may also be delineated and/or organized hierarchically according to function, geography, architectural considerations, or other factors useful or desirable in performing information management operations. A first cell may represent a geographic segment of an enterprise, such as a Chicago office, and a second cell may represent a different geographic segment, such as a New York office. Other cells may represent departments within a particular office. Where delineated by function, a first cell may perform one or more first types of information management operations (e.g., one or more first types of secondary or other copies), and a second cell may perform one or more second types of information management operations (e.g., one or more second types of secondary or other copies).

The storage manager **140** may also track information that permits it to select, designate, or otherwise identify content indices, deduplication databases, or similar databases or resources or data sets within its information management cell (or another cell) to be searched in response to certain queries. Such queries may be entered by the user via interaction with the user interface **158**. In general, the management agent **154** allows multiple information management cells to communicate with one another. For example, the information management system **100** in some cases may be one information management cell of a network of multiple cells adjacent to one another or otherwise logically related in a WAN or LAN. With this arrangement, the cells may be connected to one another through respective management agents **154**.

For instance, the management agent **154** can provide the storage manager **140** with the ability to communicate with other components within the information management system **100** (and/or other cells within a larger information management system) via network protocols and application programming interfaces (“APIs”) including, e.g., HTTP, HTTPS, FTP, REST, virtualization software APIs, cloud service provider APIs, and hosted service provider APIs. Inter-cell communication and hierarchy is described in

greater detail in e.g., U.S. Pat. Nos. 7,747,579 and 7,343,453, which are incorporated by reference herein.

Data Agents

As discussed, a variety of different types of applications **110** can operate on a given client computing device **102**, including operating systems, database applications, e-mail applications, and virtual machines, just to name a few. And, as part of the process of creating and restoring secondary copies **116**, the client computing devices **102** may be tasked with processing and preparing the primary data **112** from these various different applications **110**. Moreover, the nature of the processing/preparation can differ across clients and application types, e.g., due to inherent structural and formatting differences among applications **110**.

The one or more data agent(s) **142** are therefore advantageously configured in some embodiments to assist in the performance of information management operations based on the type of data that is being protected, at a client-specific and/or application-specific level.

The data agent **142** may be a software module or component that is generally responsible for managing, initiating, or otherwise assisting in the performance of information management operations in information management system **100**, generally as directed by storage manager **140**. For instance, the data agent **142** may take part in performing data storage operations such as the copying, archiving, migrating, and/or replicating of primary data **112** stored in the primary storage device(s) **104**. The data agent **142** may receive control information from the storage manager **140**, such as commands to transfer copies of data objects, metadata, and other payload data to the media agents **144**.

In some embodiments, a data agent **142** may be distributed between the client computing device **102** and storage manager **140** (and any other intermediate components) or may be deployed from a remote location or its functions approximated by a remote process that performs some or all of the functions of data agent **142**. In addition, a data agent **142** may perform some functions provided by a media agent **144**, or may perform other functions such as encryption and deduplication.

As indicated, each data agent **142** may be specialized for a particular application **110**, and the system can employ multiple application-specific data agents **142**, each of which may perform information management operations (e.g., perform backup, migration, and data recovery) associated with a different application **110**. For instance, different individual data agents **142** may be designed to handle Microsoft Exchange data, Lotus Notes data, Microsoft Windows file system data, Microsoft Active Directory Objects data, SQL Server data, SharePoint data, Oracle database data, SAP database data, virtual machines and/or associated data, and other types of data.

A file system data agent, for example, may handle data files and/or other file system information. If a client computing device **102** has two or more types of data, a specialized data agent **142** may be used for each data type to copy, archive, migrate, and restore the client computing device **102** data. For example, to backup, migrate, and/or restore all of the data on a Microsoft Exchange server, the client computing device **102** may use a Microsoft Exchange Mailbox data agent **142** to back up the Exchange mailboxes, a Microsoft Exchange Database data agent **142** to back up the Exchange databases, a Microsoft Exchange Public Folder data agent **142** to back up the Exchange Public Folders, and a Microsoft Windows File System data agent **142** to back up the file system of the client computing device **102**. In such embodiments, these specialized data agents **142** may be

treated as four separate data agents **142** even though they operate on the same client computing device **102**.

Other embodiments may employ one or more generic data agents **142** that can handle and process data from two or more different applications **110**, or that can handle and process multiple data types, instead of or in addition to using specialized data agents **142**. For example, one generic data agent **142** may be used to back up, migrate and restore Microsoft Exchange Mailbox data and Microsoft Exchange Database data while another generic data agent may handle Microsoft Exchange Public Folder data and Microsoft Windows File System data.

Each data agent **142** may be configured to access data and/or metadata stored in the primary storage device(s) **104** associated with the data agent **142** and process the data as appropriate. For example, during a secondary copy operation, the data agent **142** may arrange or assemble the data and metadata into one or more files having a certain format (e.g., a particular backup or archive format) before transferring the file(s) to a media agent **144** or other component. The file(s) may include a list of files or other metadata. Each data agent **142** can also assist in restoring data or metadata to primary storage devices **104** from a secondary copy **116**. For instance, the data agent **142** may operate in conjunction with the storage manager **140** and one or more of the media agents **144** to restore data from secondary storage device(s) **108**.

Media Agents

As indicated above with respect to FIG. 1A, off-loading certain responsibilities from the client computing devices **102** to intermediate components such as the media agent(s) **144** can provide a number of benefits including improved client computing device **102** operation, faster secondary copy operation performance, and enhanced scalability. In one specific example which will be discussed below in further detail, the media agent **144** can act as a local cache of copied data and/or metadata that it has stored to the secondary storage device(s) **108**, providing improved restore capabilities.

Generally speaking, a media agent **144** may be implemented as a software module that manages, coordinates, and facilitates the transmission of data, as directed by the storage manager **140**, between a client computing device **102** and one or more secondary storage devices **108**. Whereas the storage manager **140** controls the operation of the information management system **100**, the media agent **144** generally provides a portal to secondary storage devices **108**. For instance, other components in the system interact with the media agents **144** to gain access to data stored on the secondary storage devices **108**, whether it be for the purposes of reading, writing, modifying, or deleting data. Moreover, as will be described further, media agents **144** can generate and store information relating to characteristics of the stored data and/or metadata, or can generate and store other types of information that generally provides insight into the contents of the secondary storage devices **108**.

Media agents **144** can comprise separate nodes in the information management system **100** (e.g., nodes that are separate from the client computing devices **102**, storage manager **140**, and/or secondary storage devices **108**). In general, a node within the information management system **100** can be a logically and/or physically separate component, and in some cases is a component that is individually addressable or otherwise identifiable. In addition, each media agent **144** may operate on a dedicated secondary storage computing device **106** in some cases, while in other

embodiments a plurality of media agents **144** operate on the same secondary storage computing device **106**.

A media agent **144** (and corresponding media agent database **152**) may be considered to be “associated with” a particular secondary storage device **108** if that media agent **144** is capable of one or more of: routing and/or storing data to the particular secondary storage device **108**, coordinating the routing and/or storing of data to the particular secondary storage device **108**, retrieving data from the particular secondary storage device **108**, coordinating the retrieval of data from a particular secondary storage device **108**, and modifying and/or deleting data retrieved from the particular secondary storage device **108**.

While media agent(s) **144** are generally associated with one or more secondary storage devices **108**, one or more media agents **144** in certain embodiments are physically separate from the secondary storage devices **108**. For instance, the media agents **144** may operate on secondary storage computing devices **106** having different housings or packages than the secondary storage devices **108**. In one example, a media agent **144** operates on a first server computer and is in communication with a secondary storage device(s) **108** operating in a separate, rack-mounted RAID-based system.

Where the information management system **100** includes multiple media agents **144** (see, e.g., FIG. 1D), a first media agent **144** may provide failover functionality for a second, failed media agent **144**. In addition, media agents **144** can be dynamically selected for storage operations to provide load balancing. Failover and load balancing are described in greater detail below.

In operation, a media agent **144** associated with a particular secondary storage device **108** may instruct the secondary storage device **108** to perform an information management operation. For instance, a media agent **144** may instruct a tape library to use a robotic arm or other retrieval means to load or eject a certain storage media, and to subsequently archive, migrate, or retrieve data to or from that media, e.g., for the purpose of restoring the data to a client computing device **102**. As another example, a secondary storage device **108** may include an array of hard disk drives or solid state drives organized in a RAID configuration, and the media agent **144** may forward a logical unit number (LUN) and other appropriate information to the array, which uses the received information to execute the desired storage operation. The media agent **144** may communicate with a secondary storage device **108** via a suitable communications link, such as a SCSI or Fiber Channel link.

As shown, each media agent **144** may maintain an associated media agent database **152**. The media agent database **152** may be stored in a disk or other storage device (not shown) that is local to the secondary storage computing device **106** on which the media agent **144** operates. In other cases, the media agent database **152** is stored remotely from the secondary storage computing device **106**.

The media agent database **152** can include, among other things, an index **153** (see, e.g., FIG. 1C), which comprises information generated during secondary copy operations and other storage or information management operations. The index **153** provides a media agent **144** or other component with a fast and efficient mechanism for locating secondary copies **116** or other data stored in the secondary storage devices **108**. In some cases, the index **153** does not form a part of and is instead separate from the media agent database **152**.

A media agent index **153** or other data structure associated with the particular media agent **144** may include information

about the stored data. For instance, for each secondary copy **116**, the index **153** may include metadata such as a list of the data objects (e.g., files/subdirectories, database objects, mailbox objects, etc.), a path to the secondary copy **116** on the corresponding secondary storage device **108**, location information indicating where the data objects are stored in the secondary storage device **108**, when the data objects were created or modified, etc. Thus, the index **153** includes metadata associated with the secondary copies **116** that is readily available for use without having to be first retrieved from the secondary storage device **108**. In yet further embodiments, some or all of the information in index **153** may instead or additionally be stored along with the secondary copies of data in a secondary storage device **108**. In some embodiments, the secondary storage devices **108** can include sufficient information to perform a “bare metal restore”, where the operating system of a failed client computing device **102** or other restore target is automatically rebuilt as part of a restore operation.

Because the index **153** maintained in the media agent database **152** may operate as a cache, it can also be referred to as “an index cache.” In such cases, information stored in the index cache **153** typically comprises data that reflects certain particulars about storage operations that have occurred relatively recently. After some triggering event, such as after a certain period of time elapses, or the index cache **153** reaches a particular size, the index cache **153** may be copied or migrated to a secondary storage device(s) **108**. This information may need to be retrieved and uploaded back into the index cache **153** or otherwise restored to a media agent **144** to facilitate retrieval of data from the secondary storage device(s) **108**. In some embodiments, the cached information may include format or containerization information related to archives or other files stored on the storage device(s) **108**. In this manner, the index cache **153** allows for accelerated restores.

In some alternative embodiments the media agent **144** generally acts as a coordinator or facilitator of storage operations between client computing devices **102** and corresponding secondary storage devices **108**, but does not actually write the data to the secondary storage device **108**. For instance, the storage manager **140** (or the media agent **144**) may instruct a client computing device **102** and secondary storage device **108** to communicate with one another directly. In such a case the client computing device **102** transmits the data directly or via one or more intermediary components to the secondary storage device **108** according to the received instructions, and vice versa. In some such cases, the media agent **144** may still receive, process, and/or maintain metadata related to the storage operations. Moreover, in these embodiments, the payload data can flow through the media agent **144** for the purposes of populating the index cache **153** maintained in the media agent database **152**, but not for writing to the secondary storage device **108**.

The media agent **144** and/or other components such as the storage manager **140** may in some cases incorporate additional functionality, such as data classification, content indexing, deduplication, encryption, compression, and the like. Further details regarding these and other functions are described below.

Distributed, Scalable Architecture

As described, certain functions of the information management system **100** can be distributed amongst various physical and/or logical components in the system. For instance, one or more of the storage manager **140**, data agents **142**, and media agents **144** may operate on computing

devices that are physically separate from one another. This architecture can provide a number of benefits.

For instance, hardware and software design choices for each distributed component can be targeted to suit its particular function. The secondary computing devices **106** on which the media agents **144** operate can be tailored for interaction with associated secondary storage devices **108** and provide fast index cache operation, among other specific tasks. Similarly, the client computing device(s) **102** can be selected to effectively service the applications **110** thereon, in order to efficiently produce and store primary data **112**.

Moreover, in some cases, one or more of the individual components in the information management system **100** can be distributed to multiple, separate computing devices. As one example, for large file systems where the amount of data stored in the management database **146** is relatively large, the database **146** may be migrated to or otherwise reside on a specialized database server (e.g., an SQL server) separate from a server that implements the other functions of the storage manager **140**. This distributed configuration can provide added protection because the database **146** can be protected with standard database utilities (e.g., SQL log shipping or database replication) independent from other functions of the storage manager **140**. The database **146** can be efficiently replicated to a remote site for use in the event of a disaster or other data loss at the primary site. Or the database **146** can be replicated to another computing device within the same site, such as to a higher performance machine in the event that a storage manager host device can no longer service the needs of a growing information management system **100**.

The distributed architecture also provides both scalability and efficient component utilization. FIG. 1D shows an embodiment of the information management system **100** including a plurality of client computing devices **102** and associated data agents **142** as well as a plurality of secondary storage computing devices **106** and associated media agents **144**.

Additional components can be added or subtracted based on the evolving needs of the information management system **100**. For instance, depending on where bottlenecks are identified, administrators can add additional client computing devices **102**, secondary storage computing devices **106** (and corresponding media agents **144**), and/or secondary storage devices **108**. Moreover, where multiple fungible components are available, load balancing can be implemented to dynamically address identified bottlenecks. As an example, the storage manager **140** may dynamically select which media agents **144** and/or secondary storage devices **108** to use for storage operations based on a processing load analysis of the media agents **144** and/or secondary storage devices **108**, respectively.

Moreover, each client computing device **102** in some embodiments can communicate with, among other components, any of the media agents **144**, e.g., as directed by the storage manager **140**. And each media agent **144** may be able to communicate with, among other components, any of the secondary storage devices **108**, e.g., as directed by the storage manager **140**. Thus, operations can be routed to the secondary storage devices **108** in a dynamic and highly flexible manner, to provide load balancing, failover, and the like. Further examples of scalable systems capable of dynamic storage operations, and of systems capable of performing load balancing and fail over are provided in U.S. Pat. No. 7,246,207, which is incorporated by reference herein.

In alternative configurations, certain components are not distributed and may instead reside and execute on the same computing device. For example, in some embodiments, one or more data agents **142** and the storage manager **140** operate on the same client computing device **102**. In another embodiment, one or more data agents **142** and one or more media agents **144** operate on a single computing device.

Exemplary Types of Information Management Operations

In order to protect and leverage stored data, the information management system **100** can be configured to perform a variety of information management operations. As will be described, these operations can generally include secondary copy and other data movement operations, processing and data manipulation operations, analysis, reporting, and management operations. The operations described herein may be performed on any type of computing device, e.g., between two computers connected via a LAN, to a mobile client telecommunications device connected to a server via a WLAN, to any manner of client computing device coupled to a cloud storage target, etc., without limitation.

Data Movement Operations

Data movement operations according to certain embodiments are generally operations that involve the copying or migration of data (e.g., payload data) between different locations in the information management system **100** in an original/native and/or one or more different formats. For example, data movement operations can include operations in which stored data is copied, migrated, or otherwise transferred from one or more first storage devices to one or more second storage devices, such as from primary storage device(s) **104** to secondary storage device(s) **108**, from secondary storage device(s) **108** to different secondary storage device(s) **108**, from secondary storage devices **108** to primary storage devices **104**, or from primary storage device(s) **104** to different primary storage device(s) **104**.

Data movement operations can include by way of example, backup operations, archive operations, information lifecycle management operations such as hierarchical storage management operations, replication operations (e.g., continuous data replication operations), snapshot operations, deduplication or single-instancing operations, auxiliary copy operations, and the like. As will be discussed, some of these operations involve the copying, migration or other movement of data, without actually creating multiple, distinct copies. Nonetheless, some or all of these operations are referred to as “copy” operations for simplicity.

Backup Operations

A backup operation creates a copy of a version of data (e.g., one or more files or other data units) in primary data **112** at a particular point in time. Each subsequent backup copy may be maintained independently of the first. Further, a backup copy in some embodiments is generally stored in a form that is different than the native format, e.g., a backup format. This can be in contrast to the version in primary data **112** from which the backup copy is derived, and which may instead be stored in a native format of the source application(s) **110**. In various cases, backup copies can be stored in a format in which the data is compressed, encrypted, deduplicated, and/or otherwise modified from the original application format. For example, a backup copy may be stored in a backup format that facilitates compression and/or efficient long-term storage.

Backup copies can have relatively long retention periods as compared to primary data **112**, and may be stored on media with slower retrieval times than primary data **112** and certain other types of secondary copies **116**. On the other hand, backups may have relatively shorter retention periods

than some other types of secondary copies **116**, such as archive copies (described below). Backups may sometimes be stored at an offsite location.

Backup operations can include full backups, differential backups, incremental backups, “synthetic full” backups, and/or creating a “reference copy.” A full backup (or “standard full backup”) in some embodiments is generally a complete image of the data to be protected. However, because full backup copies can consume a relatively large amount of storage, it can be useful to use a full backup copy as a baseline and only store changes relative to the full backup copy for subsequent backup copies.

For instance, a differential backup operation (or cumulative incremental backup operation) tracks and stores changes that have occurred since the last full backup. Differential backups can grow quickly in size, but can provide relatively efficient restore times because a restore can be completed in some cases using only the full backup copy and the latest differential copy.

An incremental backup operation generally tracks and stores changes since the most recent backup copy of any type, which can greatly reduce storage utilization. In some cases, however, restore times can be relatively long in comparison to full or differential backups because completing a restore operation may involve accessing a full backup in addition to multiple incremental backups.

Synthetic full backups generally consolidate data without directly backing up data from the client computing device. A synthetic full backup is created from the most recent full backup (i.e., standard or synthetic) and subsequent incremental and/or differential backups. The resulting synthetic full backup is identical to what would have been created had the last backup for the subclient been a standard full backup. Unlike standard full, incremental, and differential backups, a synthetic full backup does not actually transfer data from a client computer to the backup media, because it operates as a backup consolidator. A synthetic full backup extracts the index data of each participating subclient. Using this index data and the previously backed up user data images, it builds new full backup images, one for each subclient. The new backup images consolidate the index and user data stored in the related incremental, differential, and previous full backups, in some embodiments creating an archive file at the subclient level.

Any of the above types of backup operations can be at the volume-level, file-level, or block-level. Volume level backup operations generally involve the copying of a data volume (e.g., a logical disk or partition) as a whole. In a file-level backup, the information management system **100** may generally track changes to individual files, and includes copies of files in the backup copy. In the case of a block-level backup, files are broken into constituent blocks, and changes are tracked at the block-level. Upon restore, the information management system **100** reassembles the blocks into files in a transparent fashion.

Far less data may actually be transferred and copied to the secondary storage devices **108** during a file-level copy than a volume-level copy. Likewise, a block-level copy may involve the transfer of less data than a file-level copy, resulting in faster execution times. However, restoring a relatively higher-granularity copy can result in longer restore times. For instance, when restoring a block-level copy, the process of locating constituent blocks can sometimes result in longer restore times as compared to file-level backups. Similar to backup operations, the other types of

secondary copy operations described herein can also be implemented at either the volume-level, file-level, or block-level.

For example, in some embodiments, a reference copy may comprise copy(ies) of selected objects from backed up data, typically to help organize data by keeping contextual information from multiple sources together, and/or help retain specific data for a longer period of time, such as for legal hold needs. A reference copy generally maintains data integrity, and when the data is restored, it may be viewed in the same format as the source data. In some embodiments, a reference copy is based on a specialized client, individual subclient and associated information management policies (e.g., storage policy, retention policy, etc.) that are administered within information management system **100**.

Archive Operations

Because backup operations generally involve maintaining a version of the copied data in primary data **112** and also maintaining backup copies in secondary storage device(s) **108**, they can consume significant storage capacity. To help reduce storage consumption, an archive operation according to certain embodiments creates a secondary copy **116** by both copying and removing source data. Or, seen another way, archive operations can involve moving some or all of the source data to the archive destination. Thus, data satisfying criteria for removal (e.g., data of a threshold age or size) may be removed from source storage. The source data may be primary data **112** or a secondary copy **116**, depending on the situation. As with backup copies, archive copies can be stored in a format in which the data is compressed, encrypted, deduplicated, and/or otherwise modified from the format of the original application or source copy. In addition, archive copies may be retained for relatively long periods of time (e.g., years) and, in some cases, are never deleted. Archive copies are generally retained for longer periods of time than backup copies, for example. In certain embodiments, archive copies may be made and kept for extended periods in order to meet compliance regulations.

Moreover, when primary data **112** is archived, in some cases the corresponding primary data **112** or a portion thereof is deleted when creating the archive copy. Thus, archiving can serve the purpose of freeing up space in the primary storage device(s) **104** and easing the demand on computational resources on client computing device **102**. Similarly, when a secondary copy **116** is archived, the secondary copy **116** may be deleted, and an archive copy can therefore serve the purpose of freeing up space in secondary storage device(s) **108**. In contrast, source copies often remain intact when creating backup copies. Examples of compatible data archiving operations are provided in U.S. Pat. No. 7,107,298, which is incorporated by reference herein.

Snapshot Operations

Snapshot operations can provide a relatively lightweight, efficient mechanism for protecting data. From an end-user viewpoint, a snapshot may be thought of as an “instant” image of the primary data **112** at a given point in time, and may include state and/or status information relative to an application that creates/manages the primary data **112**. In one embodiment, a snapshot may generally capture the directory structure of an object in primary data **112** such as a file or volume or other data set at a particular moment in time and may also preserve file attributes and contents. A snapshot in some cases is created relatively quickly, e.g., substantially instantly, using a minimum amount of file space, but may still function as a conventional file system backup.

A “hardware snapshot” (or “hardware-based snapshot”) operation can be a snapshot operation where a target storage device (e.g., a primary storage device **104** or a secondary storage device **108**) performs the snapshot operation in a self-contained fashion, substantially independently, using hardware, firmware and/or software operating on the storage device itself. For instance, the storage device may be capable of performing snapshot operations upon request, generally without intervention or oversight from any of the other components in the information management system **100**. In this manner, hardware snapshots can off-load other components of information management system **100** from processing involved in snapshot creation and management.

A “software snapshot” (or “software-based snapshot”) operation, on the other hand, can be a snapshot operation in which one or more other components in information management system **100** (e.g., client computing devices **102**, data agents **142**, etc.) implement a software layer that manages the snapshot operation via interaction with the target storage device. For instance, the component executing the snapshot management software layer may derive a set of pointers and/or data that represents the snapshot. The snapshot management software layer may then transmit the same to the target storage device, along with appropriate instructions for writing the snapshot.

Some types of snapshots do not actually create another physical copy of all the data as it existed at the particular point in time, but may simply create pointers that are able to map files and directories to specific memory locations (e.g., to specific disk blocks) where the data resides, as it existed at the particular point in time. For example, a snapshot copy may include a set of pointers derived from the file system or from an application. In some other cases, the snapshot may be created at the block-level, such that creation of the snapshot occurs without awareness of the file system. Each pointer points to a respective stored data block, so that collectively, the set of pointers reflect the storage location and state of the data object (e.g., file(s) or volume(s) or data set(s)) at a particular point in time when the snapshot copy was created.

An initial snapshot may use only a small amount of disk space needed to record a mapping or other data structure representing or otherwise tracking the blocks that correspond to the current state of the file system. Additional disk space is usually required only when files and directories are modified later on. Furthermore, when files are modified, typically only the pointers which map to blocks are copied, not the blocks themselves. In some embodiments, for example in the case of “copy-on-write” snapshots, when a block changes in primary storage, the block is copied to secondary storage or cached in primary storage before the block is overwritten in primary storage, and the pointer to that block is changed to reflect the new location of that block. The snapshot mapping of file system data may also be updated to reflect the changed block(s) at that particular point in time. In some other cases, a snapshot includes a full physical copy of all or substantially all of the data represented by the snapshot. Further examples of snapshot operations are provided in U.S. Pat. No. 7,529,782, which is incorporated by reference herein.

A snapshot copy in many cases can be made quickly and without significantly impacting primary computing resources because large amounts of data need not be copied or moved. In some embodiments, a snapshot may exist as a virtual file system, parallel to the actual file system. Users in some cases gain read-only access to the record of files and directories of the snapshot. By electing to restore primary

data **112** from a snapshot taken at a given point in time, users may also return the current file system to the state of the file system that existed when the snapshot was taken.

Replication Operations

Another type of secondary copy operation is a replication operation. Some types of secondary copies **116** are used to periodically capture images of primary data **112** at particular points in time (e.g., backups, archives, and snapshots). However, it can also be useful for recovery purposes to protect primary data **112** in a more continuous fashion, by replicating the primary data **112** substantially as changes occur. In some cases a replication copy can be a mirror copy, for instance, where changes made to primary data **112** are mirrored or substantially immediately copied to another location (e.g., to secondary storage device(s) **108**). By copying each write operation to the replication copy, two storage systems are kept synchronized or substantially synchronized so that they are virtually identical at approximately the same time. Where entire disk volumes are mirrored, however, mirroring can require significant amount of storage space and utilizes a large amount of processing resources.

According to some embodiments storage operations are performed on replicated data that represents a recoverable state, or “known good state” of a particular application running on the source system. For instance, in certain embodiments, known good replication copies may be viewed as copies of primary data **112**. This feature allows the system to directly access, copy, restore, backup or otherwise manipulate the replication copies as if the data were the “live” primary data **112**. This can reduce access time, storage utilization, and impact on source applications **110**, among other benefits. Based on known good state information, the information management system **100** can replicate sections of application data that represent a recoverable state rather than rote copying of blocks of data. Examples of compatible replication operations (e.g., continuous data replication) are provided in U.S. Pat. No. 7,617,262, which is incorporated by reference herein.

Deduplication/Single-Instancing Operations

Another type of data movement operation is deduplication or single-instance storage, which is useful to reduce the amount of non-primary data. For instance, some or all of the above-described secondary storage operations can involve deduplication in some fashion. New data is read, broken down into portions (e.g., sub-file level blocks, files, etc.) of a selected granularity, compared with blocks that are already in secondary storage, and only the new blocks are stored. Blocks that already exist are represented as pointers to the already stored data.

In order to streamline the comparison process, the information management system **100** may calculate and/or store signatures (e.g., hashes or cryptographically unique IDs) corresponding to the individual data blocks in a database and compare the signatures instead of comparing entire data blocks. In some cases, only a single instance of each element is stored, and deduplication operations may therefore be referred to interchangeably as “single-instancing” operations. Depending on the implementation, however, deduplication or single-instancing operations can store more than one instance of certain data blocks, but nonetheless significantly reduce data redundancy. Depending on the embodiment, deduplication blocks can be of fixed or variable length. Using variable length blocks can provide enhanced deduplication by responding to changes in the data stream, but can involve complex processing. In some cases, the information management system **100** utilizes a technique for

dynamically aligning deduplication blocks (e.g., fixed-length blocks) based on changing content in the data stream, as described in U.S. Pat. No. 8,364,652, which is incorporated by reference herein.

The information management system **100** can perform deduplication in a variety of manners at a variety of locations in the information management system **100**. For instance, in some embodiments, the information management system **100** implements “target-side” deduplication by deduplicating data (e.g., secondary copies **116**) stored in the secondary storage devices **108**. In some such cases, the media agents **144** are generally configured to manage the deduplication process. For instance, one or more of the media agents **144** maintain a corresponding deduplication database that stores deduplication information (e.g., data-block signatures). Examples of such a configuration are provided in U.S. Pat. Pub. No. 2012/0150826, which is incorporated by reference herein. Instead of or in combination with “target-side” deduplication, deduplication can also be performed on the “source-side” (or “client-side”), e.g., to reduce the amount of traffic between the media agents **144** and the client computing device(s) **102** and/or reduce redundant data stored in the primary storage devices **104**. According to various implementations, one or more of the storage devices of the target-side and/or source-side of an operation can be cloud-based storage devices. Thus, the target-side and/or source-side deduplication can be cloud-based deduplication. In particular, as discussed previously, the storage manager **140** may communicate with other components within the information management system **100** via network protocols and cloud service provider APIs to facilitate cloud-based deduplication/single instancing. Examples of such deduplication techniques are provided in U.S. Pat. Pub. No. 2012/0150818, which is incorporated by reference herein. Some other compatible deduplication/single instancing techniques are described in U.S. Pat. Pub. Nos. 2006/0224846 and 2009/0319534, which are incorporated by reference herein.

Information Lifecycle Management and Hierarchical Storage Management Operations

In some embodiments, files and other data over their lifetime move from more expensive, quick access storage to less expensive, slower access storage. Operations associated with moving data through various tiers of storage are sometimes referred to as information lifecycle management (ILM) operations.

One type of ILM operation is a hierarchical storage management (HSM) operation. A HSM operation is generally an operation for automatically moving data between classes of storage devices, such as between high-cost and low-cost storage devices. For instance, an HSM operation may involve movement of data from primary storage devices **104** to secondary storage devices **108**, or between tiers of secondary storage devices **108**. With each tier, the storage devices may be progressively relatively cheaper, have relatively slower access/restore times, etc. For example, movement of data between tiers may occur as data becomes less important over time.

In some embodiments, an HSM operation is similar to an archive operation in that creating an HSM copy may (though not always) involve deleting some of the source data, e.g., according to one or more criteria related to the source data. For example, an HSM copy may include data from primary data **112** or a secondary copy **116** that is larger than a given size threshold or older than a given age threshold and that is stored in a backup format.

Often, and unlike some types of archive copies, HSM data that is removed or aged from the source is replaced by a logical reference pointer or stub. The reference pointer or stub can be stored in the primary storage device **104** (or other source storage device, such as a secondary storage device **108**) to replace the deleted source data and to point to or otherwise indicate the new location in a secondary storage device **108**.

According to one example, files are generally moved between higher and lower cost storage depending on how often the files are accessed. When a user requests access to the HSM data that has been removed or migrated, the information management system **100** uses the stub to locate the data and may make recovery of the data appear transparent, even though the HSM data may be stored at a location different from other source data. In this manner, the data appears to the user (e.g., in file system browsing windows and the like) as if it still resides in the source location (e.g., in a primary storage device **104**). The stub may also include some metadata associated with the corresponding data, so that a file system and/or application can provide some information about the data object and/or a limited-functionality version (e.g., a preview) of the data object.

An HSM copy may be stored in a format other than the native application format (e.g., where the data is compressed, encrypted, deduplicated, and/or otherwise modified from the original native application format). In some cases, copies which involve the removal of data from source storage and the maintenance of stub or other logical reference information on source storage may be referred to generally as “on-line archive copies”. On the other hand, copies which involve the removal of data from source storage without the maintenance of stub or other logical reference information on source storage may be referred to as “off-line archive copies”. Examples of HSM and ILM techniques are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

Auxiliary Copy and Disaster Recovery Operations

An auxiliary copy is generally a copy operation in which a copy is created of an existing secondary copy **116**. For instance, an initial secondary copy **116** may be generated using or otherwise be derived from primary data **112** (or other data residing in the secondary storage subsystem **118**), whereas an auxiliary copy is generated from the initial secondary copy **116**. Auxiliary copies can be used to create additional standby copies of data and may reside on different secondary storage devices **108** than the initial secondary copies **116**. Thus, auxiliary copies can be used for recovery purposes if initial secondary copies **116** become unavailable. Exemplary compatible auxiliary copy techniques are described in further detail in U.S. Pat. No. 8,230,195, which is incorporated by reference herein.

The information management system **100** may also perform disaster recovery operations that make or retain disaster recovery copies, often as secondary, high-availability disk copies. The information management system **100** may create secondary disk copies and store the copies at disaster recovery locations using auxiliary copy or replication operations, such as continuous data replication technologies. Depending on the particular data protection goals, disaster recovery locations can be remote from the client computing devices **102** and primary storage devices **104**, remote from some or all of the secondary storage devices **108**, or both.

Data Analysis, Reporting, and Management Operations

Data analysis, reporting, and management operations can be different than data movement operations in that they do

not necessarily involve the copying, migration or other transfer of data (e.g., primary data **112** or secondary copies **116**) between different locations in the system. For instance, data analysis operations may involve processing (e.g., offline processing) or modification of already stored primary data **112** and/or secondary copies **116**. However, in some embodiments data analysis operations are performed in conjunction with data movement operations. Some data analysis operations include content indexing operations and classification operations which can be useful in leveraging the data under management to provide enhanced search and other features. Other data analysis operations such as compression and encryption can provide data reduction and security benefits, respectively.

Classification Operations/Content Indexing

In some embodiments, the information management system **100** analyzes and indexes characteristics, content, and metadata associated with the primary data **112** and/or secondary copies **116**. The content indexing can be used to identify files or other data objects having pre-defined content (e.g., user-defined keywords or phrases, other keywords/phrases that are not defined by a user, etc.), and/or metadata (e.g., email metadata such as “to”, “from”, “cc”, “bcc”, attachment name, received time, etc.).

The information management system **100** generally organizes and catalogues the results in a content index, which may be stored within the media agent database **152**, for example. The content index can also include the storage locations of (or pointer references to) the indexed data in the primary data **112** or secondary copies **116**, as appropriate. The results may also be stored, in the form of a content index database or otherwise, elsewhere in the information management system **100** (e.g., in the primary storage devices **104**, or in the secondary storage device **108**). Such index data provides the storage manager **140** or another component with an efficient mechanism for locating primary data **112** and/or secondary copies **116** of data objects that match particular criteria.

For instance, search criteria can be specified by a user through user interface **158** of the storage manager **140**. In some cases, the information management system **100** analyzes data and/or metadata in secondary copies **116** to create an “off-line” content index, without significantly impacting the performance of the client computing devices **102**. Depending on the embodiment, the system can also implement “on-line” content indexing, e.g., of primary data **112**. Examples of compatible content indexing techniques are provided in U.S. Pat. No. 8,170,995, which is incorporated by reference herein.

One or more components can be configured to scan data and/or associated metadata for classification purposes to populate a database (or other data structure) of information, which can be referred to as a “data classification database” or a “metabase”. Depending on the embodiment, the data classification database(s) can be organized in a variety of different ways, including centralization, logical sub-divisions, and/or physical sub-divisions. For instance, one or more centralized data classification databases may be associated with different subsystems or tiers within the information management system **100**. As an example, there may be a first centralized metabase associated with the primary storage subsystem **117** and a second centralized metabase associated with the secondary storage subsystem **118**. In other cases, there may be one or more metabases associated with individual components, e.g., client computing devices **102** and/or media agents **144**. In some embodiments, a data classification database (metabase) may reside as one or more

data structures within management database **146**, or may be otherwise associated with storage manager **140**.

In some cases, the metabase(s) may be included in separate database(s) and/or on separate storage device(s) from primary data **112** and/or secondary copies **116**, such that operations related to the metabase do not significantly impact performance on other components in the information management system **100**. In other cases, the metabase(s) may be stored along with primary data **112** and/or secondary copies **116**. Files or other data objects can be associated with identifiers (e.g., tag entries, etc.) in the media agent **144** (or other indices) to facilitate searches of stored data objects. Among a number of other benefits, the metabase can also allow efficient, automatic identification of files or other data objects to associate with secondary copy or other information management operations (e.g., in lieu of scanning an entire file system). Examples of compatible metabases and data classification operations are provided in U.S. Pat. Nos. 8,229,954 and 7,747,579, which are incorporated by reference herein.

Encryption Operations

The information management system **100** in some cases is configured to process data (e.g., files or other data objects, secondary copies **116**, etc.), according to an appropriate encryption algorithm (e.g., Blowfish, Advanced Encryption Standard [AES], Triple Data Encryption Standard [3-DES], etc.) to limit access and provide data security in the information management system **100**. The information management system **100** in some cases encrypts the data at the client level, such that the client computing devices **102** (e.g., the data agents **142**) encrypt the data prior to forwarding the data to other components, e.g., before sending the data to media agents **144** during a secondary copy operation. In such cases, the client computing device **102** may maintain or have access to an encryption key or passphrase for decrypting the data upon restore. Encryption can also occur when creating copies of secondary copies, e.g., when creating auxiliary copies or archive copies. In yet further embodiments, the secondary storage devices **108** can implement built-in, high performance hardware encryption.

Management and Reporting Operations

Certain embodiments leverage the integrated, ubiquitous nature of the information management system **100** to provide useful system-wide management and reporting functions. Examples of some compatible management and reporting techniques are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

Operations management can generally include monitoring and managing the health and performance of information management system **100** by, without limitation, performing error tracking, generating granular storage/performance metrics (e.g., job success/failure information, deduplication efficiency, etc.), generating storage modeling and costing information, and the like. As an example, a storage manager **140** or other component in the information management system **100** may analyze traffic patterns and suggest and/or automatically route data via a particular route to minimize congestion. In some embodiments, the system can generate predictions relating to storage operations or storage operation information. Such predictions, which may be based on a trending analysis, may predict various network operations or resource usage, such as network traffic levels, storage media use, use of bandwidth of communication links, use of media agent components, etc. Further examples of traffic analysis, trend analysis, prediction generation, and the like are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

In some configurations, a master storage manager **140** may track the status of storage operation cells in a hierarchy, such as the status of jobs, system components, system resources, and other items, by communicating with storage managers **140** (or other components) in the respective storage operation cells. Moreover, the master storage manager **140** may track the status of its associated storage operation cells and information management operations by receiving periodic status updates from the storage managers **140** (or other components) in the respective cells regarding jobs, system components, system resources, and other items. In some embodiments, a master storage manager **140** may store status information and other information regarding its associated storage operation cells and other system information in its index **150** (or other location).

The master storage manager **140** or other component may also determine whether certain storage-related criteria or other criteria are satisfied, and perform an action or trigger event (e.g., data migration) in response to the criteria being satisfied, such as where a storage threshold is met for a particular volume, or where inadequate protection exists for certain data. For instance, in some embodiments, data from one or more storage operation cells is used to dynamically and automatically mitigate recognized risks, and/or to advise users of risks or suggest actions to mitigate these risks. For example, an information management policy may specify certain requirements (e.g., that a storage device should maintain a certain amount of free space, that secondary copies should occur at a particular interval, that data should be aged and migrated to other storage after a particular period, that data on a secondary volume should always have a certain level of availability and be restorable within a given time period, that data on a secondary volume may be mirrored or otherwise migrated to a specified number of other volumes, etc.). If a risk condition or other criterion is triggered, the system may notify the user of these conditions and may suggest (or automatically implement) an action to mitigate or otherwise address the risk. For example, the system may indicate that data from a primary copy **112** should be migrated to a secondary storage device **108** to free space on the primary storage device **104**. Examples of the use of risk factors and other triggering criteria are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

In some embodiments, the system **100** may also determine whether a metric or other indication satisfies particular storage criteria and, if so, perform an action. For example, as previously described, a storage policy or other definition might indicate that a storage manager **140** should initiate a particular action if a storage metric or other indication drops below or otherwise fails to satisfy specified criteria such as a threshold of data protection. Examples of such metrics are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

In some embodiments, risk factors may be quantified into certain measurable service or risk levels for ease of comprehension. For example, certain applications and associated data may be considered to be more important by an enterprise than other data and services. Financial compliance data, for example, may be of greater importance than marketing materials, etc. Network administrators may assign priority values or “weights” to certain data and/or applications, corresponding to the relative importance. The level of compliance of storage operations specified for these applications may also be assigned a certain value. Thus, the health, impact, and overall importance of a service may be determined, such as by measuring the compliance value and

calculating the product of the priority value and the compliance value to determine the “service level” and comparing it to certain operational thresholds to determine whether it is acceptable. Further examples of the service level determination are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

The system **100** may additionally calculate data costing and data availability associated with information management operation cells according to an embodiment of the invention. For instance, data received from the cell may be used in conjunction with hardware-related information and other information about system elements to determine the cost of storage and/or the availability of particular data in the system. Exemplary information generated could include how fast a particular department is using up available storage space, how long data would take to recover over a particular system pathway from a particular secondary storage device, costs over time, etc. Moreover, in some embodiments, such information may be used to determine or predict the overall cost associated with the storage of certain information. The cost associated with hosting a certain application may be based, at least in part, on the type of media on which the data resides, for example. Storage devices may be assigned to a particular cost categories, for example. Further examples of costing techniques are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

Any of the above types of information (e.g., information related to trending, predictions, job, cell or component status, risk, service level, costing, etc.) can generally be provided to users via the user interface **158** in a single, integrated view or console (not shown). The console may support a reporting capability that allows for the generation of a variety of reports, which may be tailored to a particular aspect of information management. Report types may include: scheduling, event management, media management and data aging. Available reports may also include backup history, data aging history, auxiliary copy history, job history, library and drive, media in library, restore history, and storage policy, etc., without limitation. Such reports may be specified and created at a certain point in time as a system analysis, forecasting, or provisioning tool. Integrated reports may also be generated that illustrate storage and performance metrics, risks and storage costing information. Moreover, users may create their own reports based on specific needs.

The integrated user interface **158** can include an option to show a “virtual view” of the system that graphically depicts the various components in the system using appropriate icons. As one example, the user interface **158** may provide a graphical depiction of one or more primary storage devices **104**, the secondary storage devices **108**, data agents **142** and/or media agents **144**, and their relationship to one another in the information management system **100**. The operations management functionality can facilitate planning and decision-making. For example, in some embodiments, a user may view the status of some or all jobs as well as the status of each component of the information management system **100**. Users may then plan and make decisions based on this data. For instance, a user may view high-level information regarding storage operations for the information management system **100**, such as job status, component status, resource status (e.g., communication pathways, etc.), and other information. The user may also drill down or use other means to obtain more detailed information regarding a particular component, job, or the like. Further examples of some reporting techniques and associated interfaces providing an integrated view of an information management sys-

tem are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

The information management system **100** can also be configured to perform system-wide e-discovery operations in some embodiments. In general, e-discovery operations provide a unified collection and search capability for data in the system, such as data stored in the secondary storage devices **108** (e.g., backups, archives, or other secondary copies **116**). For example, the information management system **100** may construct and maintain a virtual repository for data stored in the information management system **100** that is integrated across source applications **110**, different storage device types, etc. According to some embodiments, e-discovery utilizes other techniques described herein, such as data classification and/or content indexing.

Information Management Policies

As indicated previously, an information management policy **148** can include a data structure or other information source that specifies a set of parameters (e.g., criteria and rules) associated with secondary copy and/or other information management operations.

One type of information management policy **148** is a storage policy. According to certain embodiments, a storage policy generally comprises a data structure or other information source that defines (or includes information sufficient to determine) a set of preferences or other criteria for performing information management operations. Storage policies can include one or more of the following items: (1) what data will be associated with the storage policy; (2) a destination to which the data will be stored; (3) datapath information specifying how the data will be communicated to the destination; (4) the type of storage operation to be performed; and (5) retention information specifying how long the data will be retained at the destination (see, e.g., FIG. 1E).

As an illustrative example, data associated with a storage policy can be logically organized into groups. In some cases, these logical groupings can be referred to as “sub-clients”. A sub-client may represent static or dynamic associations of portions of a data volume. Sub-clients may represent mutually exclusive portions. Thus, in certain embodiments, a portion of data may be given a label and the association is stored as a static entity in an index, database or other storage location. Sub-clients may also be used as an effective administrative scheme of organizing data according to data type, department within the enterprise, storage preferences, or the like. Depending on the configuration, sub-clients can correspond to files, folders, virtual machines, databases, etc. In one exemplary scenario, an administrator may find it preferable to separate e-mail data from financial data using two different sub-clients.

A storage policy can define where data is stored by specifying a target or destination storage device (or group of storage devices). For instance, where the secondary storage device **108** includes a group of disk libraries, the storage policy may specify a particular disk library for storing the sub-clients associated with the policy. As another example, where the secondary storage devices **108** include one or more tape libraries, the storage policy may specify a particular tape library for storing the sub-clients associated with the storage policy, and may also specify a drive pool and a tape pool defining a group of tape drives and a group of tapes, respectively, for use in storing the sub-client data. While information in the storage policy can be statically assigned in some cases, some or all of the information in the storage policy can also be dynamically determined based on criteria, which can be set forth in the storage policy. For

instance, based on such criteria, a particular destination storage device(s) (or other parameter of the storage policy) may be determined based on characteristics associated with the data involved in a particular storage operation, device availability (e.g., availability of a secondary storage device **108** or a media agent **144**), network status and conditions (e.g., identified bottlenecks), user credentials, and the like).

Datapath information can also be included in the storage policy. For instance, the storage policy may specify network pathways and components to utilize when moving the data to the destination storage device(s). In some embodiments, the storage policy specifies one or more media agents **144** for conveying data associated with the storage policy between the source (e.g., one or more host client computing devices **102**) and destination (e.g., a particular target secondary storage device **108**).

A storage policy can also specify the type(s) of operations associated with the storage policy, such as a backup, archive, snapshot, auxiliary copy, or the like. Retention information can specify how long the data will be kept, depending on organizational needs (e.g., a number of days, months, years, etc.)

Another type of information management policy **148** is a scheduling policy, which specifies when and how often to perform operations. Scheduling parameters may specify with what frequency (e.g., hourly, weekly, daily, event-based, etc.) or under what triggering conditions secondary copy or other information management operations will take place. Scheduling policies in some cases are associated with particular components, such as particular logical groupings of data associated with a storage policy (e.g., a sub-client), client computing device **102**, and the like. In one configuration, a separate scheduling policy is maintained for particular logical groupings of data on a client computing device **102**. The scheduling policy specifies that those logical groupings are to be moved to secondary storage devices **108** every hour according to storage policies associated with the respective sub-clients.

When adding a new client computing device **102**, administrators can manually configure information management policies **148** and/or other settings, e.g., via the user interface **158**. However, this can be an involved process resulting in delays, and it may be desirable to begin data protection operations quickly, without awaiting human intervention. Thus, in some embodiments, the information management system **100** automatically applies a default configuration to client computing device **102**. As one example, when one or more data agent(s) **142** are installed on one or more client computing devices **102**, the installation script may register the client computing device **102** with the storage manager **140**, which in turn applies the default configuration to the new client computing device **102**. In this manner, data protection operations can begin substantially immediately. The default configuration can include a default storage policy, for example, and can specify any appropriate information sufficient to begin data protection operations. This can include a type of data protection operation, scheduling information, a target secondary storage device **108**, data path information (e.g., a particular media agent **144**), and the like.

Other types of information management policies **148** are possible, including one or more audit (or security) policies. An audit policy is a set of preferences, rules and/or criteria that protect sensitive data in the information management system **100**. For example, an audit policy may define “sensitive objects” as files or objects that contain particular keywords (e.g., “confidential,” or “privileged”) and/or are associated with particular keywords (e.g., in metadata) or

particular flags (e.g., in metadata identifying a document or email as personal, confidential, etc.). An audit policy may further specify rules for handling sensitive objects. As an example, an audit policy may require that a reviewer approve the transfer of any sensitive objects to a cloud storage site, and that if approval is denied for a particular sensitive object, the sensitive object should be transferred to a local primary storage device **104** instead. To facilitate this approval, the audit policy may further specify how a secondary storage computing device **106** or other system component should notify a reviewer that a sensitive object is slated for transfer.

Another type of information management policy **148** is a provisioning policy. A provisioning policy can include a set of preferences, priorities, rules, and/or criteria that specify how client computing devices **102** (or groups thereof) may utilize system resources, such as available storage on cloud storage and/or network bandwidth. A provisioning policy specifies, for example, data quotas for particular client computing devices **102** (e.g., a number of gigabytes that can be stored monthly, quarterly or annually). The storage manager **140** or other components may enforce the provisioning policy. For instance, the media agents **144** may enforce the policy when transferring data to secondary storage devices **108**. If a client computing device **102** exceeds a quota, a budget for the client computing device **102** (or associated department) is adjusted accordingly or an alert may trigger.

While the above types of information management policies **148** have been described as separate policies, one or more of these can be generally combined into a single information management policy **148**. For instance, a storage policy may also include or otherwise be associated with one or more scheduling, audit, or provisioning policies or operational parameters thereof. Moreover, while storage policies are typically associated with moving and storing data, other policies may be associated with other types of information management operations. The following is a non-exhaustive list of items the information management policies **148** may specify:

- schedules or other timing information, e.g., specifying when and/or how often to perform information management operations;
- the type of copy **116** (e.g., type of secondary copy) and/or copy format (e.g., snapshot, backup, archive, HSM, etc.);
- a location or a class or quality of storage for storing secondary copies **116** (e.g., one or more particular secondary storage devices **108**);
- preferences regarding whether and how to encrypt, compress, deduplicate, or otherwise modify or transform secondary copies **116**;
- which system components and/or network pathways (e.g., preferred media agents **144**) should be used to perform secondary storage operations;
- resource allocation among different computing devices or other system components used in performing information management operations (e.g., bandwidth allocation, available storage capacity, etc.);
- whether and how to synchronize or otherwise distribute files or other data objects across multiple computing devices or hosted services; and
- retention information specifying the length of time primary data **112** and/or secondary copies **116** should be retained, e.g., in a particular class or tier of storage devices, or within the information management system **100**.

Policies can additionally specify or depend on a variety of historical or current criteria that may be used to determine which rules to apply to a particular data object, system component, or information management operation, such as:

- frequency with which primary data **112** or a secondary copy **116** of a data object or metadata has been or is predicted to be used, accessed, or modified;
- time-related factors (e.g., aging information such as time since the creation or modification of a data object);
- deduplication information (e.g., hashes, data blocks, deduplication block size, deduplication efficiency or other metrics);
- an estimated or historic usage or cost associated with different components (e.g., with secondary storage devices **108**);
- the identity of users, applications **110**, client computing devices **102** and/or other computing devices that created, accessed, modified, or otherwise utilized primary data **112** or secondary copies **116**;
- a relative sensitivity (e.g., confidentiality, importance) of a data object, e.g., as determined by its content and/or metadata;
- the current or historical storage capacity of various storage devices;
- the current or historical network capacity of network pathways connecting various components within the storage operation cell;
- access control lists or other security information; and
- the content of a particular data object (e.g., its textual content) or of metadata associated with the data object.

Exemplary Storage Policy and Secondary Storage Operations

FIG. 1E includes a data flow diagram depicting performance of storage operations by an embodiment of an information management system **100**, according to an exemplary storage policy **148A**. The information management system **100** includes a storage manager **140**, a client computing device **102** having a file system data agent **142A** and an email data agent **142B** operating thereon, a primary storage device **104**, two media agents **144A**, **144B**, and two secondary storage devices **108A**, **108B**: a disk library **108A** and a tape library **108B**. As shown, the primary storage device **104** includes primary data **112A**, which is associated with a logical grouping of data associated with a file system, and primary data **112B**, which is associated with a logical grouping of data associated with email. Although for simplicity the logical grouping of data associated with the file system is referred to as a file system sub-client, and the logical grouping of data associated with the email is referred to as an email sub-client, the techniques described with respect to FIG. 1E can be utilized in conjunction with data that is organized in a variety of other manners.

As indicated by the dashed box, the second media agent **144B** and the tape library **108B** are “off-site”, and may therefore be remotely located from the other components in the information management system **100** (e.g., in a different city, office building, etc.). Indeed, “off-site” may refer to a magnetic tape located in storage, which must be manually retrieved and loaded into a tape drive to be read. In this manner, information stored on the tape library **108B** may provide protection in the event of a disaster or other failure.

The file system sub-client and its associated primary data **112A** in certain embodiments generally comprise information generated by the file system and/or operating system of the client computing device **102**, and can include, for example, file system data (e.g., regular files, file tables, mount points, etc.), operating system data (e.g., registries,

event logs, etc.), and the like. The e-mail sub-client, on the other hand, and its associated primary data **112B**, include data generated by an e-mail application operating on the client computing device **102**, and can include mailbox information, folder information, emails, attachments, associated database information, and the like. As described above, the sub-clients can be logical containers, and the data included in the corresponding primary data **112A**, **1126** may or may not be stored contiguously.

The exemplary storage policy **148A** includes backup copy preferences (or rule set) **160**, disaster recovery copy preferences rule set **162**, and compliance copy preferences or rule set **164**. The backup copy rule set **160** specifies that it is associated with a file system sub-client **166** and an email sub-client **168**. Each of these sub-clients **166**, **168** are associated with the particular client computing device **102**. The backup copy rule set **160** further specifies that the backup operation will be written to the disk library **108A**, and designates a particular media agent **144A** to convey the data to the disk library **108A**. Finally, the backup copy rule set **160** specifies that backup copies created according to the rule set **160** are scheduled to be generated on an hourly basis and to be retained for 30 days. In some other embodiments, scheduling information is not included in the storage policy **148A**, and is instead specified by a separate scheduling policy.

The disaster recovery copy rule set **162** is associated with the same two sub-clients **166**, **168**. However, the disaster recovery copy rule set **162** is associated with the tape library **108B**, unlike the backup copy rule set **160**. Moreover, the disaster recovery copy rule set **162** specifies that a different media agent, namely **144B**, will be used to convey the data to the tape library **1086**. As indicated, disaster recovery copies created according to the rule set **162** will be retained for 60 days, and will be generated on a daily basis. Disaster recovery copies generated according to the disaster recovery copy rule set **162** can provide protection in the event of a disaster or other catastrophic data loss that would affect the backup copy **116A** maintained on the disk library **108A**.

The compliance copy rule set **164** is only associated with the email sub-client **168**, and not the file system sub-client **166**. Compliance copies generated according to the compliance copy rule set **164** will therefore not include primary data **112A** from the file system sub-client **166**. For instance, the organization may be under an obligation to store and maintain copies of email data for a particular period of time (e.g., 10 years) to comply with state or federal regulations, while similar regulations do not apply to the file system data. The compliance copy rule set **164** is associated with the same tape library **108B** and media agent **144B** as the disaster recovery copy rule set **162**, although a different storage device or media agent could be used in other embodiments. Finally, the compliance copy rule set **164** specifies that copies generated under the compliance copy rule set **164** will be retained for 10 years, and will be generated on a quarterly basis.

At step **1**, the storage manager **140** initiates a backup operation according to the backup copy rule set **160**. For instance, a scheduling service running on the storage manager **140** accesses scheduling information from the backup copy rule set **160** or a separate scheduling policy associated with the client computing device **102**, and initiates a backup copy operation on an hourly basis. Thus, at the scheduled time slot the storage manager **140** sends instructions to the client computing device **102** (i.e., to both data agent **142A** and data agent **142B**) to begin the backup operation.

At step **2**, the file system data agent **142A** and the email data agent **142B** operating on the client computing device **102** respond to the instructions received from the storage manager **140** by accessing and processing the primary data **112A**, **112B** involved in the copy operation, which can be found in primary storage device **104**. Because the operation is a backup copy operation, the data agent(s) **142A**, **142B** may format the data into a backup format or otherwise process the data.

At step **3**, the client computing device **102** communicates the retrieved, processed data to the first media agent **144A**, as directed by the storage manager **140**, according to the backup copy rule set **160**. In some other embodiments, the information management system **100** may implement a load-balancing, availability-based, or other appropriate algorithm to select from the available set of media agents **144A**, **144B**. Regardless of the manner the media agent **144A** is selected, the storage manager **140** may further keep a record in the storage manager database **146** of the association between the selected media agent **144A** and the client computing device **102** and/or between the selected media agent **144A** and the backup copy **116A**.

The target media agent **144A** receives the data from the client computing device **102**, and at step **4** conveys the data to the disk library **108A** to create the backup copy **116A**, again at the direction of the storage manager **140** and according to the backup copy rule set **160**. The secondary storage device **108A** can be selected in other ways. For instance, the media agent **144A** may have a dedicated association with a particular secondary storage device(s), or the storage manager **140** or media agent **144A** may select from a plurality of secondary storage devices, e.g., according to availability, using one of the techniques described in U.S. Pat. No. 7,246,207, which is incorporated by reference herein.

The media agent **144A** can also update its index **153** to include data and/or metadata related to the backup copy **116A**, such as information indicating where the backup copy **116A** resides on the disk library **108A**, data and metadata for cache retrieval, etc. The storage manager **140** may similarly update its index **150** to include information relating to the storage operation, such as information relating to the type of storage operation, a physical location associated with one or more copies created by the storage operation, the time the storage operation was performed, status information relating to the storage operation, the components involved in the storage operation, and the like. In some cases, the storage manager **140** may update its index **150** to include some or all of the information stored in the index **153** of the media agent **144A**. After the 30 day retention period expires, the storage manager **140** instructs the media agent **144A** to delete the backup copy **116A** from the disk library **108A**. Indexes **150** and/or **153** are updated accordingly.

At step **5**, the storage manager **140** initiates the creation of a disaster recovery copy **1166** according to the disaster recovery copy rule set **162**.

At step **6**, illustratively based on the instructions received from the storage manager **140** at step **5**, the specified media agent **144B** retrieves the most recent backup copy **116A** from the disk library **108A**.

At step **7**, again at the direction of the storage manager **140** and as specified in the disaster recovery copy rule set **162**, the media agent **144B** uses the retrieved data to create a disaster recovery copy **116B** on the tape library **108B**. In some cases, the disaster recovery copy **116B** is a direct, mirror copy of the backup copy **116A**, and remains in the backup format. In other embodiments, the disaster recovery

copy 1166 may be generated in some other manner, such as by using the primary data 112A, 112B from the primary storage device 104 as source data. The disaster recovery copy operation is initiated once a day and the disaster recovery copies 1166 are deleted after 60 days; indexes are updated accordingly when/after each information management operation is executed/completed.

At step 8, the storage manager 140 initiates the creation of a compliance copy 116C, according to the compliance copy rule set 164. For instance, the storage manager 140 instructs the media agent 144B to create the compliance copy 116C on the tape library 108B at step 9, as specified in the compliance copy rule set 164. In the example, the compliance copy 116C is generated using the disaster recovery copy 116B. In other embodiments, the compliance copy 116C is instead generated using either the primary data 112B corresponding to the email sub-client or using the backup copy 116A from the disk library 108A as source data. As specified, in the illustrated example, compliance copies 116C are created quarterly, and are deleted after ten years, and indexes are kept up-to-date accordingly.

While not shown in FIG. 1E, at some later point in time, a restore operation can be initiated involving one or more of the secondary copies 116A, 1166, 116C. As one example, a user may manually initiate a restore of the backup copy 116A by interacting with the user interface 158 of the storage manager 140. The storage manager 140 then accesses data in its index 150 (and/or the respective storage policy 148A) associated with the selected backup copy 116A to identify the appropriate media agent 144A and/or secondary storage device 108A.

In other cases, a media agent may be selected for use in the restore operation based on a load balancing algorithm, an availability based algorithm, or other criteria. The selected media agent 144A retrieves the data from the disk library 108A. For instance, the media agent 144A may access its index 153 to identify a location of the backup copy 116A on the disk library 108A, or may access location information residing on the disk 108A itself.

When the backup copy 116A was recently created or accessed, the media agent 144A accesses a cached version of the backup copy 116A residing in the index 153, without having to access the disk library 108A for some or all of the data. Once it has retrieved the backup copy 116A, the media agent 144A communicates the data to the source client computing device 102. Upon receipt, the file system data agent 142A and the email data agent 142B may unpackage (e.g., restore from a backup format to the native application format) the data in the backup copy 116A and restore the unpackaged data to the primary storage device 104.

Exemplary Applications of Storage Policies

The storage manager 140 may permit a user to specify aspects of the storage policy 148A. For example, the storage policy can be modified to include information governance policies to define how data should be managed in order to comply with a certain regulation or business objective. The various policies may be stored, for example, in the management database 146. An information governance policy may comprise a classification policy, which is described herein. An information governance policy may align with one or more compliance tasks that are imposed by regulations or business requirements. Examples of information governance policies might include a Sarbanes-Oxley policy, a HIPAA policy, an electronic discovery (E-Discovery) policy, and so on.

Information governance policies allow administrators to obtain different perspectives on all of an organization's

online and offline data, without the need for a dedicated data silo created solely for each different viewpoint. As described previously, the data storage systems herein build a centralized index that reflects the contents of a distributed data set that spans numerous clients and storage devices, including both primary and secondary copies, and online and offline copies. An organization may apply multiple information governance policies in a top-down manner over that unified data set and indexing schema in order to permit an organization to view and manipulate the single data set through different lenses, each of which is adapted to a particular compliance or business goal. Thus, for example, by applying an E-discovery policy and a Sarbanes-Oxley policy, two different groups of users in an organization can conduct two very different analyses of the same underlying physical set of data copies, which may be distributed throughout the organization and information management system.

A classification policy defines a taxonomy of classification terms or tags relevant to a compliance task and/or business objective. A classification policy may also associate a defined tag with a classification rule. A classification rule defines a particular combination of criteria, such as users who have created, accessed or modified a document or data object; file or application types; content or metadata keywords; clients or storage locations; dates of data creation and/or access; review status or other status within a workflow (e.g., reviewed or un-reviewed); modification times or types of modifications; and/or any other data attributes in any combination, without limitation. A classification rule may also be defined using other classification tags in the taxonomy. The various criteria used to define a classification rule may be combined in any suitable fashion, for example, via Boolean operators, to define a complex classification rule. As an example, an E-discovery classification policy might define a classification tag "privileged" that is associated with documents or data objects that (1) were created or modified by legal department staff, or (2) were sent to or received from outside counsel via email, or (3) contain one of the following keywords: "privileged" or "attorney" or "counsel", or other like terms.

One specific type of classification tag, which may be added to an index at the time of indexing, is an entity tag. An entity tag may be, for example, any content that matches a defined data mask format. Examples of entity tags might include, e.g., social security numbers (e.g., any numerical content matching the formatting mask XXX-XX-XXXX), credit card numbers (e.g., content having a 13-16 digit string of numbers), SKU numbers, product numbers, etc.

A user may define a classification policy by indicating criteria, parameters or descriptors of the policy via a graphical user interface, such as a form or page with fields to be filled in, pull-down menus or entries allowing one or more of several options to be selected, buttons, sliders, hypertext links or other known user interface tools for receiving user input, etc. For example, a user may define certain entity tags, such as a particular product number or project ID code that is relevant in the organization. In some implementations, the classification policy can be implemented using cloud-based techniques. For example, the storage devices may be cloud storage devices, and the storage manager 140 may execute cloud service provider API over a network to classify data stored on cloud storage devices.

Exemplary Secondary Copy Formatting

The formatting and structure of secondary copies 116 can vary, depending on the embodiment. In some cases, secondary copies 116 are formatted as a series of logical data units or "chunks" (e.g., 512 MB, 1 GB, 2 GB, 4 GB, or 8 GB

chunks). This can facilitate efficient communication and writing to secondary storage devices **108**, e.g., according to resource availability. For example, a single secondary copy **116** may be written on a chunk-by-chunk basis to a single secondary storage device **108** or across multiple secondary storage devices **108**. In some cases, users can select different chunk sizes, e.g., to improve throughput to tape storage devices.

Generally, each chunk can include a header and a payload. The payload can include files (or other data units) or subsets thereof included in the chunk, whereas the chunk header generally includes metadata relating to the chunk, some or all of which may be derived from the payload. For example, during a secondary copy operation, the media agent **144**, storage manager **140**, or other component may divide the associated files into chunks and generate headers for each chunk by processing the constituent files. The headers can include a variety of information such as file identifier(s), volume(s), offset(s), or other information associated with the payload data items, a chunk sequence number, etc. Importantly, in addition to being stored with the secondary copy **116** on the secondary storage device **108**, the chunk headers can also be stored to the index **153** of the associated media agent(s) **144** and/or the index **150**. This is useful in some cases for providing faster processing of secondary copies **116** during restores or other operations. In some cases, once a chunk is successfully transferred to a secondary storage device **108**, the secondary storage device **108** returns an indication of receipt, e.g., to the media agent **144** and/or storage manager **140**, which may update their respective indexes **153**, **150** accordingly. During restore, chunks may be processed (e.g., by the media agent **144**) according to the information in the chunk header to reassemble the files.

Data can also be communicated within the information management system **100** in data channels that connect the client computing devices **102** to the secondary storage devices **108**. These data channels can be referred to as “data streams”, and multiple data streams can be employed to parallelize an information management operation, improving data transfer rate, among providing other advantages. Example data formatting techniques including techniques involving data streaming, chunking, and the use of other data structures in creating copies (e.g., secondary copies) are described in U.S. Pat. Nos. 7,315,923 and 8,156,086, and 8,578,120, each of which is incorporated by reference herein.

FIGS. **1F** and **1G** are diagrams of example data streams **170** and **171**, respectively, which may be employed for performing data storage operations. Referring to FIG. **1F**, the data agent **142** forms the data stream **170** from the data associated with a client computing device **102** (e.g., primary data **112**). The data stream **170** is composed of multiple pairs of stream header **172** and stream data (or stream payload) **174**. The data streams **170** and **171** shown in the illustrated example are for a single-instanced storage operation, and a stream payload **174** therefore may include both single-instance (“SI”) data and/or non-SI data. A stream header **172** includes metadata about the stream payload **174**. This metadata may include, for example, a length of the stream payload **174**, an indication of whether the stream payload **174** is encrypted, an indication of whether the stream payload **174** is compressed, an archive file identifier (ID), an indication of whether the stream payload **174** is single instanceable, and an indication of whether the stream payload **174** is a start of a block of data.

Referring to FIG. **1G**, the data stream **171** has the stream header **172** and stream payload **174** aligned into multiple

data blocks. In this example, the data blocks are of size 64 KB. The first two stream header **172** and stream payload **174** pairs comprise a first data block of size 64 KB. The first stream header **172** indicates that the length of the succeeding stream payload **174** is 63 KB and that it is the start of a data block. The next stream header **172** indicates that the succeeding stream payload **174** has a length of 1 KB and that it is not the start of a new data block. Immediately following stream payload **174** is a pair comprising an identifier header **176** and identifier data **178**. The identifier header **176** includes an indication that the succeeding identifier data **178** includes the identifier for the immediately previous data block. The identifier data **178** includes the identifier that the data agent **142** generated for the data block. The data stream **171** also includes other stream header **172** and stream payload **174** pairs, which may be for SI data and/or for non-SI data.

FIG. **1H** is a diagram illustrating the data structures **180** that may be used to store blocks of SI data and non-SI data on the storage device (e.g., secondary storage device **108**). According to certain embodiments, the data structures **180** do not form part of a native file system of the storage device. The data structures **180** include one or more volume folders **182**, one or more chunk folders **184/185** within the volume folder **182**, and multiple files within the chunk folder **184**. Each chunk folder **184/185** includes a metadata file **186/187**, a metadata index file **188/189**, one or more container files **190/191/193**, and a container index file **192/194**. The metadata file **186/187** stores non-SI data blocks as well as links to SI data blocks stored in container files. The metadata index file **188/189** stores an index to the data in the metadata file **186/187**. The container files **190/191/193** store SI data blocks. The container index file **192/194** stores an index to the container files **190/191/193**. Among other things, the container index file **192/194** stores an indication of whether a corresponding block in a container file **190/191/193** is referred to by a link in a metadata file **186/187**. For example, data block **B2** in the container file **190** is referred to by a link in the metadata file **187** in the chunk folder **185**. Accordingly, the corresponding index entry in the container index file **192** indicates that the data block **B2** in the container file **190** is referred to. As another example, data block **B1** in the container file **191** is referred to by a link in the metadata file **187**, and so the corresponding index entry in the container index file **192** indicates that this data block is referred to.

As an example, the data structures **180** illustrated in FIG. **1H** may have been created as a result of two storage operations involving two client computing devices **102**. For example, a first storage operation on a first client computing device **102** could result in the creation of the first chunk folder **184**, and a second storage operation on a second client computing device **102** could result in the creation of the second chunk folder **185**. The container files **190/191** in the first chunk folder **184** would contain the blocks of SI data of the first client computing device **102**. If the two client computing devices **102** have substantially similar data, the second storage operation on the data of the second client computing device **102** would result in the media agent **144** storing primarily links to the data blocks of the first client computing device **102** that are already stored in the container files **190/191**. Accordingly, while a first storage operation may result in storing nearly all of the data subject to the storage operation, subsequent storage operations involving similar data may result in substantial data storage space savings, because links to already stored data blocks can be stored instead of additional instances of data blocks.

If the operating system of the secondary storage computing device **106** on which the media agent **144** operates supports sparse files, then when the media agent **144** creates container files **190/191/193**, it can create them as sparse files. A sparse file is type of file that may include empty space (e.g., a sparse file may have real data within it, such as at the beginning of the file and/or at the end of the file, but may also have empty space in it that is not storing actual data, such as a contiguous range of bytes all having a value of zero). Having the container files **190/191/193** be sparse files allows the media agent **144** to free up space in the container files **190/191/193** when blocks of data in the container files **190/191/193** no longer need to be stored on the storage devices. In some examples, the media agent **144** creates a new container file **190/191/193** when a container file **190/191/193** either includes 100 blocks of data or when the size of the container file **190** exceeds 50 MB. In other examples, the media agent **144** creates a new container file **190/191/193** when a container file **190/191/193** satisfies other criteria (e.g., it contains from approximately 100 to approximately 1000 blocks or when its size exceeds approximately 50 MB to 1 GB).

In some cases, a file on which a storage operation is performed may comprise a large number of data blocks. For example, a 100 MB file may comprise 400 data blocks of size 256 KB. If such a file is to be stored, its data blocks may span more than one container file, or even more than one chunk folder. As another example, a database file of 20 GB may comprise over 40,000 data blocks of size 512 KB. If such a database file is to be stored, its data blocks will likely span multiple container files, multiple chunk folders, and potentially multiple volume folders. Restoring such files may require accessing multiple container files, chunk folders, and/or volume folders to obtain the requisite data blocks.

Synchronizing Selected Portions of Data in a Storage Management System

FIG. 2 is a block diagram illustrating some salient portions of a storage management system **200** for partially synchronizing primary data files based on synchronizing portions thereof via auto-restore operations from backup data, according to an illustrative embodiment of the present invention. Storage management system **200**, which may be an embodiment of an enhanced information management system, illustratively comprises: primary storage device **104-1**, which is in communication with client computing device **202-1**, and comprises a primary data file **212A**, which comprises a non-synchronization portion **213A** and a synchronization portion **214A**; primary storage device **104-2**, which is in communication with client computing device **202-2**, and comprises a primary data file **212B**, which comprises a non-synchronization portion **213B** and a synchronization portion **214B**; secondary storage computing device **106** comprising media agent **144**; secondary storage device **108**, which comprises secondary copy **216**, and which is in communication with secondary storage computing device **106**; client computing device **202-1** comprising enhanced data agent **242A**; client computing device **202-2** comprising enhanced data agent **242B**; enhanced storage manager **240**; and console **291** in communication with enhanced storage manager **240**. The components may be interconnected as shown by the solid arrows. The physical communications infrastructure required to support these connections is well known in the art and may be any suitable electronic communications infrastructure, such as that

described in regard to communication pathways **114** above. Components **104**, **106**, **108**, and **144** were described in detail above.

FIG. 2 also illustrates (in the balloon emanating from console **291**) an initial administrative operation in which synchronization portions **214** (e.g., **214A** and **214B**) are identified as the subject portions of their respective parent files **212** to cross-synchronize according to the illustrative embodiment. These selected portions may be referred to herein as “synchronization portions” or “sync portions.” Because only selected sync portions synchronize, the respective parent data files comprising these selected sync portions may be said to “partially synchronize.” This partial-synchronization approach means that parent data files are not fully synchronized, which is in marked contrast to prior art approaches that keep files synchronized in their entirety. The present approach may be particularly advantageous when parent data files are very large and synchronization would use up a great deal of bandwidth and processing resources. Instead, partial synchronization allows for smaller amounts of data to traverse the various networks and components of system **200**. One exemplary use case for partial synchronization is for source code in a large base; when a number of different developers are working off a large base of source code that may change frequently, it may be advantageous for them to synchronize certain relevant portions to the source code base and to each other (but not the entirety of the base). The illustrative embodiment enables such a scenario to occur.

FIG. 2 also depicts an illustrative partial-synchronization operation, according to the illustrative embodiment, using the dotted arrows. The “BACKUP” arrow represents a backup operation in which the primary data that is to be synchronized, e.g., sync portion **214A**, is backed up to secondary storage by generating a secondary copy **216** which is stored to secondary storage device **108**. The “RESTORE” arrow represents a restore operation that may immediately and automatically follow the backup; accordingly, secondary copy **216** may be restored to sync portion **214B**, which is in primary storage. The result is that primary data files **212A** and **212B** partially synchronize to capture the change to the sync portion **214A** by synchronizing their respective sync portions **214A** and **214B**. More details on these operations are provided in subsequent figures.

Although not shown in the present figure, synchronization between sync portions **214A** and **214B** may also occur in the opposite direction, i.e., detecting a change in sync portion **214B**, backing up sync portion **214B**, and restoring from backup to portion **214A**. This scenario is described in more detail in reference to FIG. 5. Thus, partial synchronization according to the illustrative embodiment may initiate at any synchronization portion that undergoes a change and may then propagate to other one or more associated sync portions; this mutual process may be referred to herein as “cross-synchronization.”

In regard to non-synchronization portions **213A** and **213B**, there is no limitation to the differences therebetween, e.g., respective type, size, format, associated application, storage device, etc. Moreover, there is no limitation in regard to the location and/or placement of a sync portion **214** within a respective parent file **212** and/or relative to another sync portion within another parent file **212**.

Client computing devices **202** (e.g., **202-1** and **202-2**) are analogous to client computing device **102** described in further detail above and may comprise additional features for operating in storage management system **200**, e.g., a respective enhanced data agent **242**. Like client computing

device 102, client computing device 202 also comprises one or more applications 110, which execute thereon (not shown in the present figure). Each application 110 may generate primary data, such as primary data files 212A and 212B.

Primary data files 212 (e.g., 212A and 212B) comprise primary or “production” data generated by a client computing device 202, e.g., executing an application 110 such as a database management system or a file manager, etc. without limitation. Primary data is described in more detail in regard to element 112 above. Illustratively, a primary data file 212 comprises at least one non-synchronization portion 213 and at least one synchronization portion 214. For simplicity, primary data files 212A and 212B are illustrated with only one sync portion 214 and only one non-sync portion 213, but no such limitation is required for the present invention.

Non-synchronization portions 213 (e.g., portion 213A and portion 213B) comprise data that is not synchronized by components of the illustrative embodiment. Thus, non-sync portion 213A and non-sync portion 213B may diverge or converge over time whether in content, size, type of data, data format, etc. For example, non-sync portion 213A may represent a smaller fraction of parent primary file 212A than non-sync portion 213B relative to parent primary file 212B.

Synchronization portions 214 (e.g., 214A and 214B) comprise data is synchronized across the parent primary data files 212, according to the illustrative embodiment. Sync portion 214A may represent a smaller fraction of parent primary file 212A than sync portion 214B relative to parent primary file 212B. More details on how storage management system manages the cross-synchronization of portions 214 are provided below and in the accompanying figures.

Secondary copy 216 comprises a copy of primary data, i.e., it is not itself primary “production” data. Secondary copies are described in detail in regard to element 116 above. In the example depicted in the present figure, secondary copy 216 comprises a copy of a sync portion 214, e.g., sync portion 214A, and is generated in a partial-sync operation; furthermore, secondary copy 216 is restored to sync portion 214B, thus keeping sync portion 214A and 214B synchronized with each other. Sync portions 214 may cross-synchronize bi-directionally. An alternative embodiment that also enables bi-directional synchronization among sync portions 214 may utilize a first secondary copy 216 for backing up sync portion 214A and a second secondary copy 166 for backing up sync portion 214B. A corresponding restore operation may follow the backup operation, thus keeping sync portions 214A and 214B synchronized with each other.

Enhanced storage manager 240 is analogous to storage manager 140 described in detail above, and additionally comprises enhanced functionality for operating in storage management system 200 (e.g., partial synchronization manager module 347 described in FIG. 3). More details are given below and in regard to the accompanying figures.

Enhanced data agents 242 (e.g., 242A, 242B) are analogous to data agent 142 described in detail above and further comprise additional functionality for operating in storage management system 200 (e.g., a respective change monitor module 324 described in FIG. 3). More details are given below and in regard to the accompanying figures.

Console 291 is a component well known in the art, which, according to the illustrative embodiment, displays to a human user an interface provided by enhanced storage manager 240. In some alternative embodiments, the user interface may provide web-console access to storage manager 240.

FIG. 3 depicts some salient details of system 200 according to the illustrative embodiment. In addition to the com-

ponents depicted in the preceding figure, storage management system 200 further comprises: operating-system change log 302-1 in client computing device 202-1; operating-system change log 302-2 in client computing device 202-2; change monitor 342A in enhanced data agent 242A; change monitor 342B in enhanced data agent 242B; synchronization database 346 in enhanced storage manager 240; partial-synchronization manager 347 and jobs agent 356 in enhanced storage manager 240. The arrows in FIG. 3 depict a number of non-exclusive logical interconnections between some of the illustrative components of system 200. More details about the operations supported by these illustrative interconnections are described below and in regard to the accompanying figures.

Operating-system change log 302-1, which resides and operates in client computing device 202-1, is well known in the art. Operating-system change log 302-1, which is a utility provided by the operating system in client computing device 202-1, tracks changes to data associated with client computing device 202-1 that are occurring as a result of application(s) or any computing operations executing on client computing device 202-1. Thus, according to the illustrative configuration of system 200, operating-system change log 302-1 may track changes in data stored in primary storage device 104-1, which data includes file 212A. For example, when an application writes a changed block to file 212A, operating-system change log 302-1 creates a corresponding entry indicating that the respective data block has changed.

Operating-system change log 302-2 in client computing device 202-2 is analogous to change log 302-1, and operates in reference to computing device 202-2 and its associated data in primary storage device 104-2, including data file 212B.

Change monitor 342 (e.g., 342A, 342B) is a functional component of enhanced data agent 242, and may be implemented as executable software and/or firmware, which executes on the underlying client computing device 202 that hosts the enhanced data agent 242. When it executes according to the illustrative embodiment, change monitor 342 is largely responsible for detecting that a change has occurred in a monitored sync portion of a data file, e.g., 214A, 214B, etc. Accordingly, change monitor 342 may receive metadata from enhanced storage manager 240, identifying which sync portions 214 to monitor relative to the client computing device 202 on which change monitor 342 executes. When it detects a change to a sync portion 214, change monitor 342 may notify storage manager 240. More details regarding the functionality of change monitor 342 are given below and in the accompanying figures.

Change monitor 342 is shown herein as a distinct component within enhanced data agent 342 to ease understanding of the present disclosure, however, alternative embodiments are also possible within the scope of the present invention. Change monitor 342 may be layered on existing data agent code, or may be a logical construct whose functionality is distributed through one or more other functional modules of the data agent—and in any combination thereof.

Synchronization database 346 resides in (or is associated with) enhanced storage manager 240. Synchronization database 346 may comprise metadata for managing partial synchronization in system 200. Examples of metadata stored in synchronization database 346 may include identification markers (e.g., block identifiers, offsets, etc.) for sync portions 214 in respective data files 212; associations between a sync portion 214 in a given primary data file (e.g., 214A)

and other sync portions **214** in other primary data files (e.g., **214B**) that are to be synchronized; and other metadata without limitation. For example, synchronization database **346** may store partial-synchronization selections supplied by administrator(s) and/or user(s) (see, e.g., depicted balloon). For example, an administrator may have selected sync portions **214A** and **214B** for partial synchronization; accordingly an association between them as well as their respective identities/locations may be stored in synchronization database **346**.

Notably, there may be a plurality of different groupings of associated sync portions that are to be kept mutually synchronized in system **200**, such that each grouping may comprise sync portions that are distinct from those in other groupings. Synchronization database **346** may be part of management database **146**, which is described in detail elsewhere herein. In some embodiments, synchronization database **346** may be configured separately from management database **346**.

Partial-synchronization manager **347** is a functional component of enhanced storage manager **240**, and may be implemented as executable software and/or firmware, which executes on the underlying computing device that hosts storage manager **240**. When it executes according to the illustrative embodiment, manager **347** is largely responsible for managing partial-synchronization operations in storage management system **200**. For example, manager **347** may perform one or more of the following operations, without limitation: presenting a user interface to console **291** to enable administrators/users to select and associate sync portions **214** that they wish to keep synchronized via partial-synchronization operations in system **200**; generate and store appropriate metadata for said selections and associations and store the metadata to synchronization database **346** and retrieve metadata therefrom; transmit appropriate metadata to each change monitor **342** involved in monitoring a given sync portion **214**; receive notice of change in a given sync portion **214** (e.g., from a change monitor **342**) and in response manage a partial-sync operation relative to said changed sync portion **214**; track partial-sync operations and maintain job status. Managing a partial-sync operation may include automatically launching a backup of a changed sync portion **214** (hereinafter the “source sync portion”) and automatically restoring the backup sync data to an associated sync portion (hereinafter the “target sync portion”). Notably, the partial-sync operations are managed so that no administrator/user intervention is required after selecting and associating the sync portions, i.e., partial-sync operations are automatic and managed largely by storage manager **240** (e.g., using partial sync manager **347** and jobs agent **356**). More details are given below and in the accompanying figures.

Partial-synchronization manager **347** is shown herein as a distinct component to ease understanding of the present disclosure, however, alternative embodiments are also possible within the scope of the present invention. Manager **347** may be embodied as a unified module within storage manager **240**, layered on existing storage manager code, or may be a logical construct whose functionality is distributed through one or more other functional modules of the storage manager, such as management agent **154** and/or jobs agent **156**—and in any combination thereof. In some alternative embodiments, partial-synchronization manager **347** may execute on another computing component that is physically distinct from storage manager **240**, such as on a dedicated

server that operates in conjunction with storage manager **240**, but which may or may not be a logical component of storage manager **240**.

Jobs agent **356** is a functional component of enhanced storage manager **240** and is analogous to jobs agent **156** described in detail above and additionally comprises functionality for operating in system **200**, e.g., for interoperating with partial-sync manager **347**. Functionality for interoperating with partial-sync manager **347** may include enhanced features, message sets, data structures, etc., without limitation.

FIG. **4** depicts some salient operations of a method **400** for partially synchronizing primary data files based on synchronizing portions thereof via auto-restore operations from backup data (e.g., when source portion **214A** changes) according to an illustrative embodiment of the present invention. FIG. **4** depicts operations performed by enhanced data agent **242A**, enhanced storage manager **240**, and enhanced data agent **242B** in storage management system **200**, without limitation.

At block **401**, enhanced storage manager **240** may receive identifications of sync portions **214** that are supplied by a user/administrator (e.g., via console **291**). For example, enhanced storage manager **240** may receive user-supplied selections of sync portions **214A** and **214B**, in data files **212A** and **212B**, respectively.

At block **403**, enhanced storage manager **240** (e.g., using partial-sync manager **347**) may generate metadata comprising partial-sync relationships based on the user/administrator input received in the preceding operation. For example, enhanced storage manager **240** (e.g., using partial-sync manager **347**) may generate an association between sync portions **214A** and **214B** for ongoing partial-sync operations in system **200**.

At block **405**, enhanced storage manager **240** (e.g., using partial-sync manager **347**) may store to sync database **346** the metadata generated in the preceding operation and/or any other relevant metadata for managing partial-sync operations in storage management system **200**.

At block **407**, enhanced storage manager **240** (e.g., using partial-sync manager **347**) may transmit relevant partial-sync metadata to appropriate data agents **242**. For example, storage manager **240** may transmit the identity of sync portion **214A** and an association thereof with sync portion **214B** to enhanced data agent **242A**; and may transmit the identity of sync portion **214B** and an association thereof with sync portion **214A** to enhanced data agent **242B**.

At block **409**, enhanced data agent **242A** (e.g., using change monitor **342A**) may receive the partial-sync metadata transmitted by storage manager **240** in the preceding block. This may include the identity(ies) of sync portions **214** (e.g., **214A**) including the identities of constituent data blocks, their respective location(s) on storage device **104-1**, etc. to be monitored for changes by enhanced data agent **242A**.

At block **410**, which is analogous to block **409**, enhanced data agent **242B** (e.g., using change monitor **342B**) receives partial-sync metadata from storage manager **240**, analogous to block **409**. This may include the identity(ies) of sync portions **214** (e.g., **214B**) including the identities of constituent data blocks, their respective location(s) on storage device **104-2**, etc., to be monitored for changes by enhanced data agent **242B**.

At block **411**, enhanced data agent **242A** (e.g., using change monitor **342A**) monitors whether changes are occurring in its assigned sync portion **214A**. This operation may comprise enhanced data agent **242A** periodically checking

operating-system change log **302-1** for any changes logged in reference to one or more data blocks in sync portion **214A**.

At block **412**, which is analogous to block **411**, enhanced data agent **242B** (e.g., using change monitor **342B**) monitors whether changes are occurring in its assigned sync portion **214B**. This operation may comprise enhanced data agent **242B** periodically checking operating-system change log **302-2** for any changes logged in reference to one or more data blocks in sync portion **214B**.

At block **413**, enhanced data agent **242A** (e.g., using change monitor **342A**) may detect one or more changed data blocks in its assigned sync portions **214** (e.g., in **214A**). Sync portion **214A** is designated a source sync portion. Based on detecting said changes in the source sync portion, enhanced data agent **242A** may notify storage manager **240** of the change and/or the identities of any changed blocks.

At block **415**, enhanced storage manager **240** (e.g., using partial-sync manager **347**), may receive notice of changed blocks in a source sync portion (e.g., **214A**) from a data agent **242** (e.g., **242A**), which may be referred to as the “source data agent.” Based on the received notice, storage manager **240** (e.g., using partial-sync manager **347**), may launch and manage a partial-sync operation to all affected target sync portions **214**, as described in more detail in a subsequent figure. The result of executing block **415** is that change(s) in a source sync portion may be propagated throughout system **200** to one or more associated target sync portions, so that the respective parent data files **212** are partially synchronized relative to those associated sync portions. The present illustrative example depicts only one target sync portion **214B**, but there is no limitation to the number of associated sync portions **214** that may be managed by storage manager **240** in reference to any number of enhanced data agents **242**.

At block **417**, enhanced data agent **242A** (e.g., using change monitor **342A**) may execute a backup operation in reference to the changed source sync portion **214A**. Thus, data agent **242A** may receive an instruction from enhanced storage manager **240** to generate a secondary copy of sync portion **214A** and interoperate with media agent **144** to store the secondary copy to a secondary storage device under the management of the media agent **144** (e.g., secondary copy **216**). Enhanced data agent **242A** may communicate status and/or completion message(s) to storage manager **240** when the secondary copy **216** has been successfully stored to secondary storage.

At block **419**, enhanced data agent **242B** (e.g., using change monitor **342B**) may execute a restore operation of the secondary copy **216** to target sync portion **214B**. Thus data agent **242B** may receive an instruction from enhanced storage manager **240** to interoperate with media agent **144** (or another media agent in communication with media agent **144**) to retrieve newly created secondary copy **216** from secondary storage and transfer (e.g., copy, move, etc.) secondary copy **216** to target sync portion **214B**.

At block **421**, enhanced storage manager **240** may receive job completion status message(s) from the one or more target data agents **242** (e.g., **242B**) and may enter updated metadata to sync database **346** and/or to management database **146**, indicating that the partial-sync operation begun at block **415** is complete, and the change detected in source sync portion **214A** has been propagated successfully to target sync portion **214B** (and to any other associated sync portions **214**).

From block **421**, control may pass back to the monitoring blocks, e.g., **411**, **412**, to detect further changes in sync portions **214** throughout system **200**. A scenario in which a change is detected in sync portion **214B** is depicted in the next figure.

FIG. **5** depicts some salient operations of a method **500** for partially synchronizing primary data files based on synchronizing portions thereof via auto-restore operations from backup data (e.g., when source portion **214B** changes) according to the illustrative embodiment. FIG. **5** depicts operations performed by enhanced data agent **242A**, enhanced storage manager **240**, and enhanced data agent **242B** in storage management system **200**, without limitation. Blocks **401**, **403**, **405**, **407**, **409-412**, and **415** were described in a preceding figure. Block **415** is additionally described in more detail in a subsequent figure.

At block **513**, which follows block **412**, enhanced data agent **242B** (e.g., using change monitor **342B**) may detect one or more changed data blocks in its assigned sync portions **214** (e.g., in **214B**). Sync portion **214B** is designated a source sync portion. Based on detecting said changes in the source sync portion, enhanced data agent **242B** may notify storage manager **240** of the change and/or the identities of any changed blocks. Control passes to block **415**.

At block **517**, which follows block **415**, enhanced data agent **242B** (e.g., using change monitor **342B**) may execute a backup operation in reference to the changed source sync portion **214B**. Thus, data agent **242B** may receive an instruction from enhanced storage manager **240** to generate a secondary copy of sync portion **214B** and interoperate with media agent **144** to store the secondary copy to a secondary storage device under the management of the media agent **144** (e.g., secondary copy **216**). Enhanced data agent **242B** may communicate status and/or completion message(s) to storage manager **240** when the secondary copy **216** has been successfully stored to secondary storage.

At block **519**, which follows block **415** and block **517**, enhanced data agent **242A** (e.g., using change monitor **342A**) may execute a restore operation of the secondary copy **216** to target sync portion **214A**. Thus data agent **242A** may receive an instruction from enhanced storage manager **240** to interoperate with media agent **144** (or another media agent in communication with media agent **144**) to retrieve newly created secondary copy **216** from secondary storage and transfer (e.g., copy, move, etc.) secondary copy **216** to target sync portion **214A**.

At block **521**, which follows block **519** and the completion of the restore operation, enhanced storage manager **240** may receive job completion status message(s) from the one or more target data agents **242** (e.g., **242A**) and may enter updated metadata to sync database **346** and/or to management database **146**, indicating that the partial-sync operation begun at block **415** is complete, and the change detected in source sync portion **214B** has been propagated successfully to target sync portion **214A** (and to any other associated sync portions **214**).

From block **521**, control may pass back to the monitoring blocks, e.g., **411**, **412**, to detect further changes in sync portions **214** throughout system **200**. A scenario in which a change is detected in sync portion **214A** is depicted in the preceding figure.

FIG. **6** depicts some salient sub-operations of block **415** in method **400** and method **500**. In general, block **415** is executed by enhanced storage manager **240** (e.g., using partial-sync manager **347** and jobs agent **356**) and is directed

to launch and management of a partial-sync operation to all affected target data files based on a change detected in a source sync portion **214**.

At block **601**, enhanced storage manager **240** may receive, from a source data agent **242**, notice of change in a sync portion **214** (designated source sync portion) that is monitored by the source data agent **242**. In method **400**, storage manager **240** illustratively receives notice of change in sync portion **214A** from data agent **242A**. In method **500**, storage manager **240** illustratively receives notice of change in sync portion **214B** from data agent **242B**.

At block **603**, enhanced storage manager **240** may manage a backup operation based on the notice of change; thus, storage manager **240** may instruct the source data agent **242** and a media agent **144** to generate a secondary copy of the (changed) source sync portion **214** and store the secondary copy (e.g., **216**) to a secondary storage device (e.g., **108**); this backup operation may further include appropriate updates to the media agent index (e.g., **153**) and storage manager metadata (e.g., **356**) so that the status, location, and other characteristics of the backup copy **216** may be tracked. In contrast to prior art backup operations, which are directed at whole files, whether a full backup or an incremental one, the present backup operation is based on a distinctly defined and identified sync portion, which is specifically identified, monitored, and tracked in relation with corresponding associated sync portions in other data files in system **200**. Thus, the present backup is directed only at the sync portion **214**, which experienced a change as detected by the monitoring data agent **242**.

At block **605**, enhanced storage manager **240** may prepare for a restore operation based on partial-sync relationships metadata retrieved from synchronization database **346**. Accordingly, storage manager **240** may identify other sync portions **214** (which are designated target sync portions) associated with the source sync portion and may further identify their respective monitoring data agents **242** (which are designated target data agents).

At block **607**, enhanced storage manager **240** may instruct each target data agent **242** and media agent **144**, which stored secondary copy **216** to secondary storage (and may possibly instruct other appropriate media agents deployed elsewhere in system **200**), to restore secondary copy **216** to one or more target sync portions **214**, thereby propagating the change(s) in the source sync portion to all associated target sync portions. The restore operation(s) may further include appropriate updates to the respective media agent index (e.g., **153**) and storage manager metadata (e.g., **356**) so that the status, location, and other characteristics of the restored data may be tracked. In contrast to prior art restore operations, which are directed at whole files, the present restore operation is based on a distinctly defined and identified sync portion, which is specifically identified, monitored, and tracked in relation with corresponding associated sync portions in other data files in system **200**. Thus, the present restore is directed only at the sync portion **214**, which experienced a change as detected by the monitoring source data agent **242**.

In regard to the components, blocks, operations and/or sub-operations described in reference to FIGS. **2-6**, other embodiments are possible within the scope of the present invention, such that the above-recited components, steps, blocks, operations, and/or messages/requests/queries/instructions are differently arranged, sequenced, sub-divided, organized, and/or combined. In some embodiments, a different component may initiate or execute a given operation.

Example Embodiments

A number of example embodiments are recited below, at least in part referencing an enhanced data agent.

An illustrative method for partially synchronizing primary data files based on synchronizing portions thereof in a storage management system may comprise: receiving, by a first data agent executing on a first client computing device in a storage management system, first metadata that identifies a first portion of a first primary data file, wherein the first primary data file comprises data generated by the first client computing device, and wherein the first primary data file is associated with the first data agent; detecting, by the first data agent, at least one change in the first portion of the first primary data file; transmitting, by the first data agent to a storage manager, notice of the detected at least one change in the first portion of the first primary data file; executing, by the first data agent, at least part of a backup operation of the first portion of the first primary data file, wherein the backup operation generates a secondary copy of the first portion of the primary data file; executing, by a second data agent, at least part of a restore operation of the secondary copy to a second portion of a second primary data file, wherein the second primary data file is associated with the second data agent, which executes on a second client computing device in the storage management system; and wherein the first primary data file and the second primary data file are partially synchronized relative to the first portion and the second portion, respectively, based on the backup operation of the first portion and further based on the restore operation of the secondary copy of the first portion to the second portion of the second primary data file.

The above-recited method wherein the first primary data file and the second primary data file are partially synchronized relative to the first portion and the second portion, respectively, regardless of any changes to other portions of the first primary data file that fall outside the first portion, and also regardless of any changes to other portions of the second primary data file that fall outside the second portion. The above-recited method wherein the first data agent executes the at least part of the backup operation based on an instruction received from the storage manager. The above-recited method wherein the second data agent executes the at least part of the restore operation based on an instruction received from the storage manager. The above-recited method wherein the detecting is based at least in part on an operating-system change log residing on the first client computing device. The above-recited method wherein the detecting is based at least in part on the first data agent monitoring an operating-system change log, which resides on the first client computing device, for changes to one or more data blocks in the first portion of the first primary data file, wherein the monitoring is based on the received first metadata.

The above-recited method may further comprise: monitoring, by the first data agent, based on the received first metadata, the first portion of the first primary data file to detect whether one or more changes occur in the first portion. The above-recited method may further comprise: monitoring, by the first data agent, based on the received first metadata, an operating-system change log, which resides on the first client computing device, for changes to one or more data blocks in the first portion of the first primary data file. The above-recited method wherein the detecting of the at least one change in the first portion of the first primary data file is based at least in part on monitoring of the first portion by the first data agent, wherein the monitoring is based on the received first metadata. The above-recited method

wherein the detecting of the at least one change in the first portion of the first primary data file is based at least in part on monitoring of the first portion by the first data agent, wherein the monitoring is based at least in part on an operating-system change log residing on the first client computing device.

The above-recited method may further comprise: monitoring, by the first data agent, based on the received first metadata, the first portion of the first primary data file to detect whether one or more changes occur in the first portion; monitoring, by the second data agent, based on received second metadata that identifies the second portion of the second primary data file, to detect whether one or more changes occur in the second portion; and when a change is detected in one of the first portion and the second portion, synchronizing the first portion and the second portion to include the detected change, thereby partially synchronizing the first primary data file and the second data file, respectively; wherein the synchronizing comprises (i) backing up the changed one of the first portion and the second portion to a secondary copy, and (ii) restoring the secondary copy to the other one of the first portion and the second portion that was previously unchanged. The above-recited method wherein the first portion of the first primary data file and the second portion of the second primary data file are selected by a user of the storage management system. The above-recited method may further comprise: detecting, by the second data agent, at least one change in the second portion of the second primary data file; transmitting, by the second data agent to the storage manager, notice of the detected at least one change in the second portion of the second primary data file; based on a third instruction received from the storage manager, executing, by the second data agent, at least part of a backup operation of the second portion of the second primary data file, which results in a secondary copy of the second portion of the primary data file; based on a fourth instruction received from the storage manager, executing, by the first data agent, at least part of a restore operation of the secondary copy to the first portion of the first primary data file; and wherein the first primary data file and the second primary data file are partially synchronized relative to the first portion and the second portion, respectively, based on the backup operation of the second portion and further based on the restore operation of the secondary copy of the second portion to the first portion of the first primary data file.

Another illustrative method for partially synchronizing primary data files based on synchronizing portions thereof in a storage management system may comprise: receiving, by a first data agent executing on a first client computing device in the storage management system, first metadata that identifies a first portion of a first primary data file, wherein the first primary data file comprises data generated by the first client computing device, and wherein the first primary data file is associated with the first data agent; receiving, by a second data agent executing on a second client computing device in the storage management system, second metadata that identifies a second portion of a second primary data file, wherein the second primary data file comprises data generated by the second client computing device, and wherein the second primary data file is associated with the second data agent; monitoring the first portion of the first primary data file, by the first data agent, based on the received first metadata, to detect whether one or more changes occur in the first portion; monitoring the second portion of the second primary data file, by the second data agent, based on the second received metadata, to detect whether one or more

changes occur in the second portion; and when a change is detected in one of the first portion of the first primary data file and the second portion of the second primary data file, by a respective associated data agent, synchronizing the first portion and the second portion to include the detected change, thereby partially-synchronizing the first primary data file and the second primary data file, respectively; wherein the synchronizing comprises (i) backing up, at least in part by the detecting associated data agent, the changed one of the first portion and the second portion to a secondary copy, and (ii) restoring the secondary copy to the other one of the first portion and the second portion. The above-recited method wherein the synchronizing is based on instructions received by the respective first and second data agent from a storage manager, and further wherein the first metadata and the second metadata also are received from the storage manager. The above-recited method may further comprise: when the change is detected in one of the first portion of the first primary data file and the second portion of the second primary data file, by a respective associated data agent, transmitting notice of the detected change by the respective associated data agent to a storage manager; and wherein the synchronizing is based on instructions received by the respective first and second data agent from the storage manager. The above-recited method wherein the monitoring of the first portion by the first media agent comprises monitoring an operating-system change log, which resides on the first client computing device, for changes to one or more data blocks in the first portion of the first primary data file. The above-recited method wherein the change is detected in one of the first portion of the first primary data file and the second portion of the second primary data file, by a respective associated data agent that monitors an operating-system change log, wherein the operating-system change log resides on the same client computing device that executes the respective associated data agent that detects the change.

Another exemplary method for partially synchronizing primary data files based on synchronizing portions thereof in a storage management system may comprise: receiving, by a first data agent executing on a first client computing device in the storage management system, first metadata that identifies a first portion of a first primary data file, wherein the first primary data file comprises data generated by the first client computing device, and wherein the first primary data file is associated with the first data agent; detecting, by the first data agent, at least one change in the first portion of the first primary data file; backing up, by the first data agent in conjunction with a media agent, the first portion of the first primary data file to a secondary copy of the first portion of the primary data file, wherein the secondary copy is stored to a secondary storage device in communication with the media agent; restoring, by a second data agent in conjunction with the media agent, the secondary copy of the first portion of the primary data file to a second portion of a second primary data file, wherein the second primary data file is associated with the second data agent, which executes on a second client computing device in the storage management system; and wherein the first primary data file and the second primary data file are partially synchronized relative to the first portion and the second portion, respectively, based on the backing up of the first portion and further based on the restoring of the secondary copy of the first portion to the second portion of the second primary data file. The above-recited method wherein the backing up is executed based on instructions received by the first data agent and the media agent from a storage manager that manages the

storage management system. The above-recited method wherein the backing up is executed based on instructions received by the first data agent and the media agent from a storage manager that manages the storage management system, and further wherein the first metadata is also received from the storage manager. The above-recited method wherein the restoring is executed based on instructions received by the second data agent and the media agent from a storage manager that manages the storage management system. The above-recited method wherein the restoring is executed based on instructions received by the second data agent and the media agent from a storage manager that manages the storage management system, and further wherein the second metadata is also received from the storage manager.

A number of example embodiments are recited below, at least in part referencing an enhanced storage management system and/or enhanced storage manager.

An exemplary storage management system for partial synchronization of primary data files based on synchronization of portions thereof may comprise: a first client computing device comprising a first data agent; a first primary storage device in communication with the first client computing device, wherein the first primary storage device comprises a first primary data file generated by the first client computing device, and wherein the first primary data file comprises a first portion; a second client computing device comprising a second data agent; a second primary storage device in communication with the second client computing device, wherein the second primary storage device comprises a second primary data file generated by the second client computing device, and wherein the second primary data file comprises a second portion; a secondary storage computing device comprising a media agent; a secondary storage device in communication with the secondary storage computing device; a storage manager in communication with the first client computing device, the second client computing device, and the secondary storage computing device, wherein the storage manager comprises metadata for partial-synchronization operations in the storage management system; wherein the first data agent is configured to: receive, from the storage manager, first metadata that identifies the first portion of the first primary data file, detect at least one change in the first portion of the first primary data file, and back up, in conjunction with the media agent, the first portion of the first primary data file to a secondary copy stored to the secondary storage device; and wherein the second data agent is configured to: restore, in conjunction with the media agent, the secondary copy to the second portion of the second primary data file, thereby synchronizing the first portion of the first primary data file and the second portion of the second primary data file. The above-recited storage management system wherein the storage manager is configured to: instruct the first data agent and the media agent to back up the first portion of the first primary data file to the secondary copy, and instruct the second data agent and the media agent, after the backup of the first portion is complete, to restore the secondary copy to the second portion of the second primary data file. The above-recited storage management system wherein the storage manager is further configured to instruct the first data agent and the media agent to back up the first portion of the first primary data file based on a notice of a detected change to the first portion of the first primary data file received by the storage manager from the first data agent. The above-recited storage management system wherein the storage manager is further configured to: store first metadata that

identifies the first portion of the first primary data file, second metadata that identifies the second portion of the second primary data file, and an association between the first portion and the second portion, transmit the first metadata to the first data agent, and transmit the second metadata to the second data agent. The above-recited storage management system wherein the storage manager is configured to manage a partial-synchronization operation for the first primary data file and the second primary data file, wherein the storage manager is further configured to: instruct the first data agent and the media agent to back up the first portion of the first primary data file to the secondary copy, and instruct the second data agent and the media agent, after the backup of the first portion is complete, to restore the secondary copy to the second portion of the second primary data file.

The above-recited storage management system wherein the first data agent is configured to detect at least one change in the first portion of the first primary data file based on an operating-system change log in the first client computing device. The above-recited storage management system wherein the first primary data file and the second primary data file are partially synchronized relative to the first portion and the second portion, respectively, regardless of any changes to other portions of the first primary data file that fall outside the first portion, and also regardless of any changes to other portions of the second primary data file that fall outside the second portion. The above-recited storage management system wherein the second data agent is further configured to: detect at least one change in the second portion of the second primary data file, and back up, in conjunction with the media agent, the second portion of the second primary data file to the secondary copy stored to the secondary storage device; and wherein the first data agent is configured to: restore, in conjunction with the media agent, the secondary copy of the second portion to the first portion of the first primary data file, thereby synchronizing the first portion of the first primary data file and the second portion of the second primary data file.

Another exemplary storage management system for partial synchronization of primary data files based on synchronization of portions thereof may comprise: a first client computing device comprising a first data agent; a first primary storage device in communication with the first client computing device, wherein the first primary storage device comprises a first primary data file generated by the first client computing device, and wherein the first primary data file comprises a first portion; a second client computing device comprising a second data agent; a second primary storage device in communication with the second client computing device, wherein the second primary storage device comprises a second primary data file generated by the second client computing device, and wherein the second primary data file comprises a second portion; a secondary storage computing device comprising a media agent; a secondary storage device in communication with the secondary storage computing device; a storage manager in communication with the first client computing device, the second client computing device, and the secondary storage computing device, wherein the storage manager comprises metadata for partial-synchronization operations in the storage management system; wherein the storage manager is configured to manage a partial-synchronization operation for the first primary data file and the second primary data file, wherein the storage manager is configured to: based on a notice of a detected change to the first portion of the first primary data file received from the first data agent, instruct the first data agent and the media agent to back up the first

portion of the first primary data file to a secondary copy stored to the secondary storage device, and instruct the second data agent and the media agent, after the backup of the first portion is complete, to restore the secondary copy to the second portion of the second primary data file. The above-recited storage management system wherein the instruction to restore the secondary copy of the first portion of the first primary data file to the second portion of the second primary data file is based on an association between the first portion and the second portion. The above-recited storage management system wherein the storage manager is further configured to: store first metadata that identifies the first portion of the first primary data file, second metadata that identifies the second portion of the second primary data file, and an association between the first portion and the second portion, transmit the first metadata to the first data agent, and transmit the second metadata to the second data agent. The above-recited storage management system wherein the storage manager is further configured to: store first metadata that identifies the first portion of the first primary data file, second metadata that identifies the second portion of the second primary data file, and an association between the first portion and the second portion, transmit the first metadata to the first data agent, and transmit the second metadata to the second data agent; wherein the instruction to restore the secondary copy of the first portion of the first primary data file to the second portion of the second primary data file is based on the association between the first portion and the second portion.

The above-recited storage management system wherein the storage manager is further configured to: store first metadata that identifies the first portion of the first primary data file, second metadata that identifies the second portion of the second primary data file, and an association between the first portion and the second portion, transmit the first metadata to the first data agent, instruct the first data agent to monitor changes to the first portion of the first primary data file, transmit the second metadata to the second data agent, and instruct the second data agent to monitor changes to the second portion of the second primary data file. The above-recited storage management system wherein after the partial-synchronization operation managed by the storage manager, the first primary data file and the second primary data file are partially synchronized relative to the first portion and the second portion, respectively, regardless of any changes to other portions of the first primary data file that fall outside the first portion, and also regardless of any changes to other portions of the second primary data file that fall outside the second portion.

An exemplary method for partially synchronizing primary data files based on synchronizing portions thereof in a storage management system may comprise: storing, by a storage manager, first metadata that identifies a first portion of a first primary data file, second metadata that identifies a second portion of a second primary data file, and an association between the first portion and the second portion; managing, by the storage manager, a partial-synchronization operation for the first primary data file and the second primary data file relative to the first portion and the second portion, respectively, based at least in part on the association between the first portion and the second portion, wherein the managing of the partial-synchronization operation is based on a notice of a detected change to the first portion of the first primary data file received from a first data agent executing in the storage management system, and wherein the managing comprises: instructing the first data agent and a media agent executing in the storage management system to back

up the first portion of the first primary data file to a secondary copy, wherein the secondary copy is stored to a secondary storage device in communication with the media agent, and after the backup of the first portion is complete, instructing the media agent and a second data agent that is associated with the second primary data file to restore the secondary copy to the second portion of the second primary data file. The above-recited method may further comprise: transmitting, by the storage manager, the first metadata to the first data agent; and transmitting, by the storage manager, the second metadata to the second data agent. The above-recited method may further comprise: transmitting, by the storage manager, the first metadata to the first data agent; instructing, by the storage manager, the first data agent to monitor changes to the first portion of the first primary data file; transmitting, by the storage manager, the second metadata to the second data agent; and instructing, by the storage manager, the second data agent to monitor changes to the second portion of the second primary data file. The above-recited method wherein after the partial-synchronization operation managed by the storage manager, the first primary data file and the second primary data file are partially synchronized relative to the first portion and the second portion, respectively, regardless of any changes to other portions of the first primary data file that fall outside the first portion, and also regardless of any changes to other portions of the second primary data file that fall outside the second portion. The above-recited method may further comprise: receiving, by the first data agent executing on a first client computing device in the storage management system, the first metadata that identifies the first portion of the first primary data file, wherein the first primary data file comprises data generated by the first client computing device; detecting, by the first data agent, at least one change in the first portion of the first primary data file; transmitting, by the first data agent to a storage manager, notice of the detected at least one change in the first portion of the first primary data file; executing, by the first data agent, at least part of the backup operation of the first portion of the first primary data file to the secondary copy of the first portion of the primary data file. The above-recited method may further comprise: executing, by the second data agent, which executes on a second client computing device in the storage management system, at least part of the restore operation of the secondary copy to the second portion of a second primary data file.

Terminology

Conditional language, such as, among others, “can,” “could,” “might,” or “may,” unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey that certain embodiments include, while other embodiments do not include, certain features, elements and/or steps. Thus, such conditional language is not generally intended to imply that features, elements and/or steps are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without user input or prompting, whether these features, elements and/or steps are included or are to be performed in any particular embodiment.

Unless the context clearly requires otherwise, throughout the description and the claims, the words “comprise,” “comprising,” and the like are to be construed in an inclusive sense, as opposed to an exclusive or exhaustive sense; that is to say, in the sense of “including, but not limited to.” As used herein, the terms “connected,” “coupled,” or any variant thereof means any connection or coupling, either direct or indirect, between two or more elements; the coupling or

connection between the elements can be physical, logical, or a combination thereof. Additionally, the words “herein,” “above,” “below,” and words of similar import, when used in this application, refer to this application as a whole and not to any particular portions of this application. Where the context permits, words in the above Detailed Description using the singular or plural number may also include the plural or singular number respectively. The word “or” in reference to a list of two or more items, covers all of the following interpretations of the word: any one of the items in the list, all of the items in the list, and any combination of the items in the list. Likewise the term “and/or” in reference to a list of two or more items, covers all of the following interpretations of the word: any one of the items in the list, all of the items in the list, and any combination of the items in the list.

Depending on the embodiment, certain operations, acts, events, or functions of any of the algorithms described herein can be performed in a different sequence, can be added, merged, or left out altogether (e.g., not all are necessary for the practice of the algorithms). Moreover, in certain embodiments, operations, acts, functions, or events can be performed concurrently, e.g., through multi-threaded processing, interrupt processing, or multiple processors or processor cores or on other parallel architectures, rather than sequentially.

Systems and modules described herein may comprise software, firmware, hardware, or any combination(s) of software, firmware, or hardware suitable for the purposes described herein. Software and other modules may reside and execute on servers, workstations, personal computers, computerized tablets, PDAs, and other computing devices suitable for the purposes described herein. Software and other modules may be accessible via local memory, via a network, via a browser, or via other means suitable for the purposes described herein. Data structures described herein may comprise computer files, variables, programming arrays, programming structures, or any electronic information storage schemes or methods, or any combinations thereof, suitable for the purposes described herein. User interface elements described herein may comprise elements from graphical user interfaces, interactive voice response, command line interfaces, and other suitable interfaces.

Further, the processing of the various components of the illustrated systems can be distributed across multiple machines, networks, and other computing resources. In addition, two or more components of a system can be combined into fewer components. Various components of the illustrated systems can be implemented in one or more virtual machines, rather than in dedicated computer hardware systems and/or computing devices. Likewise, the data repositories shown can represent physical and/or logical data storage, including, for example, storage area networks or other distributed storage systems. Moreover, in some embodiments the connections between the components shown represent possible paths of data flow, rather than actual connections between hardware. While some examples of possible connections are shown, any of the subset of the components shown can communicate with any other subset of components in various implementations.

Embodiments are also described above with reference to flow chart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products. Each block of the flow chart illustrations and/or block diagrams, and combinations of blocks in the flow chart illustrations and/or block diagrams, may be implemented by computer program instructions. Such instructions may be provided to

a processor of a general purpose computer, special purpose computer, specially-equipped computer (e.g., comprising a high-performance database server, a graphics subsystem, etc.) or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor(s) of the computer or other programmable data processing apparatus, create means for implementing the acts specified in the flow chart and/or block diagram block or blocks.

These computer program instructions may also be stored in a non-transitory computer-readable memory that can direct a computer or other programmable data processing apparatus to operate in a particular manner, such that the instructions stored in the computer-readable memory produce an article of manufacture including instruction means which implement the acts specified in the flow chart and/or block diagram block or blocks. The computer program instructions may also be loaded onto a computing device or other programmable data processing apparatus to cause a series of operations to be performed on the computing device or other programmable apparatus to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide steps for implementing the acts specified in the flow chart and/or block diagram block or blocks.

Any patents and applications and other references noted above, including any that may be listed in accompanying filing papers, are incorporated herein by reference. Aspects of the invention can be modified, if necessary, to employ the systems, functions, and concepts of the various references described above to provide yet further implementations of the invention.

These and other changes can be made to the invention in light of the above Detailed Description. While the above description describes certain examples of the invention, and describes the best mode contemplated, no matter how detailed the above appears in text, the invention can be practiced in many ways. Details of the system may vary considerably in its specific implementation, while still being encompassed by the invention disclosed herein. As noted above, particular terminology used when describing certain features or aspects of the invention should not be taken to imply that the terminology is being redefined herein to be restricted to any specific characteristics, features, or aspects of the invention with which that terminology is associated. In general, the terms used in the following claims should not be construed to limit the invention to the specific examples disclosed in the specification, unless the above Detailed Description section explicitly defines such terms. Accordingly, the actual scope of the invention encompasses not only the disclosed examples, but also all equivalent ways of practicing or implementing the invention under the claims.

To reduce the number of claims, certain aspects of the invention are presented below in certain claim forms, but the applicant contemplates the various aspects of the invention in any number of claim forms. For example, while only one aspect of the invention is recited as a means-plus-function claim under 35 U.S.C. sec. 112(f) (AIA), other aspects may likewise be embodied as a means-plus-function claim, or in other forms, such as being embodied in a computer-readable medium. Any claims intended to be treated under 35 U.S.C. § 112(f) will begin with the words “means for”, but use of the term “for” in any other context is not intended to invoke treatment under 35 U.S.C. § 112(f). Accordingly, the applicant reserves the right to pursue additional claims after filing this application, in either this application or in a continuing application.

61

What is claimed is:

1. A storage management system for synchronizing portions but not entireties of primary data files based on backup and restore operations, the storage management system comprising:

a first computing device comprising one or more processors and computer-readable memory;

a second computing device comprising one or more processors and computer-readable memory;

wherein the first computing device is configured to:

detect at least one change in a first portion of a first primary data file generated by the first computing device, and

back up the first portion of the first primary data file to a secondary copy of the first portion, which is in a backup format not accessible as primary data to the first computing device;

wherein the second computing device is configured to:

restore the secondary copy of the first portion to a second portion of a second primary data file generated by the second computing device,

thereby transforming the secondary copy of the first portion from the backup format to primary data that is accessible to the second computing device within the second primary data file;

a storage manager in communication with the first computing device and the second computing device, wherein the storage manager is configured to:

based on receiving from the first computing device notice of the at least one change in the first portion of the first primary data file, (a) instruct the first computing device to back up the first portion of the first primary data file to the secondary copy of the first portion, and (b) instruct the second computing device to restore the secondary copy of the first portion to the second portion of the second primary data file; and

wherein after the secondary copy of the first portion is restored to the second portion of the second primary data file, the first portion of the first primary data file and the second portion of the second primary data file are synchronized with each other without the rest of the first primary data file and the rest of the second primary data file being respectively synchronized with each other.

2. The storage management system of claim 1, wherein after the secondary copy of the first portion is restored to the second portion of the second primary data file, the first primary data file and the second primary data file are partially synchronized with each other, based on the back up and restore operations in the storage management system.

3. The storage management system of claim 1, wherein the storage manager comprises metadata for partially synchronizing the first primary data file and the second primary data file in the storage management system relative to the first portion of the first primary data file and the second portion of the second primary data file, respectively.

4. The storage management system of claim 1, wherein the first computing device is further configured to:

receive, from the storage manager, first metadata that identifies the first portion of the first primary data file and second metadata that identifies the second portion of the second primary data file; and

wherein the storage manager comprises metadata for partially synchronizing the first primary data file and the second primary data file in the storage management system, based on synchronizing the first portion but not

62

other portions of the first primary data file and the second portion but not other portions of the second primary data file.

5. The storage management system of claim 1, wherein the first computing device comprises a first data agent associated with a first application that executes on the first computing device and wherein the first application accesses the first primary data file.

6. The storage management system of claim 1, wherein the first computing device comprises a first data agent associated with a first application that executes on the first computing device and wherein the first application accesses the first primary data file; and

wherein the second computing device comprise a second data agent associated with a second application that executes on the second computing device and wherein the second application accesses the second primary data file, including the second portion comprising the restored secondary copy of the first portion of the first primary data file.

7. The storage management system of claim 1, wherein the second computing device is further configured to:

detect at least one change in the second portion of the second primary data file, and

back up the second portion of the second primary data file to a secondary copy of the second portion, which is in a backup format not accessible as primary data to the second computing device; and

wherein the first computing device is further configured to:

restore the secondary copy of the second portion to the first portion of the first primary data file,

thereby transforming the secondary copy of the second portion from its respective backup format to primary data that is accessible to the first computing device within the first primary data file.

8. The storage management system of claim 1, wherein the first computing device is further configured to:

receive from the storage manager metadata identifying the first portion of the first primary data file enabling the first computing device to detect the at least one change in the first portion of the first primary data file.

9. The storage management system of claim 1, wherein the first computing device is further configured to:

receive from the storage manager metadata identifying the first portion of the first primary data file enabling the first computing device to detect the at least one change in the first portion of the first primary data file; and

wherein the storage manager comprises metadata for synchronizing the first portion but not other portions of the first primary data file with the second portion but not other portions of the second primary data file.

10. A method for synchronizing portions but not entireties of primary data files based on backup and restore operations in a storage management system, the method comprising:

detecting, by a first computing device comprising one or more processors and computer memory, at least one change in a first portion of a first primary data file generated by the first computing device;

based on the detecting, backing up, by the first computing device, the first portion of the first primary data file to a secondary copy of the first portion, which is in a backup format not accessible as primary data to the first computing device;

restoring, by a second computing device comprising one or more processors and computer memory, the second-

63

ary copy of the first portion to a second portion of a second primary data file generated by the second computing device,

thereby transforming the secondary copy of the first portion from the backup format to primary data that is accessible to the second computing device within the second primary data file;

based on receiving from the first computing device notice of the at least one change in the first portion of the first primary data file, instructing, by a storage manager comprising one or more processors and computer memory: (a) the first computing device to back up the first portion of the first primary data file to the secondary copy of the first portion, and (b) the second computing device to restore the secondary copy of the first portion to the second portion of the second primary data file; and

wherein after the secondary copy of the first portion is restored to the second portion of the second primary data file, the first portion of the first primary data file and the second portion of the second primary data file are synchronized with each other without the rest of the first primary data file and the rest of the second primary data file being respectively synchronized with each other.

11. The method of claim **10**, wherein the storage manager comprises metadata for partially synchronizing the first primary data file and the second primary data file in the storage management system relative to the first portion of the first primary data file and the second portion of the second primary data file, respectively.

12. The method of claim **10**, further comprising: receiving by the first computing device from the storage manager, first metadata that identifies the first portion of the first primary data file and second metadata that identifies the second portion of the second primary data file; and

wherein the storage manager comprises metadata for partially synchronizing the first primary data file and the second primary data file in the storage management system, based on synchronizing the first portion but not other portions of the first primary data file and the second portion but not other portions of the second primary data file.

13. The method of claim **10**, wherein a first application and a first data agent associated therewith execute on the first computing device and wherein the first application accesses the first primary data file; and

wherein a second application and a second data agent associated therewith execute on the second computing device and wherein the second application accesses the second primary data file, including the second portion comprising the restored secondary copy of the first portion of the first primary data file.

14. A storage management system for synchronizing portions but not entireties of primary data files based on backup and restore operations, the storage management system comprising:

a first computing device comprising one or more processors and computer-readable memory;

a first data storage device comprising a first primary data file generated by the first computing device;

a second computing device comprising one or more processors and computer-readable memory;

a second data storage device comprising a second primary data file generated by the second computing device;

64

wherein a first data agent on the first computing device is configured to:

detect at least one change in a first portion of the first primary data file, and

in conjunction with a media agent, perform a backup operation to create a secondary copy of the first portion of the first primary data file, wherein the secondary copy is in a backup format not accessible as primary data to the first computing device;

wherein a second data agent on the second computing device is configured to:

in conjunction with the media agent, restore the secondary copy of the first portion to a second portion of the second primary data file,

thereby transforming the secondary copy of the first portion from the backup format to primary data that is accessible to the second computing device within the second primary data file;

a storage manager in communication with the first data agent, the second data agent, and the media agent,

wherein the storage manager triggers the backup operation to create the secondary copy of the first portion of the first primary data file and also triggers the restore operation to the second portion of the second primary data file; and

wherein after the restore operation, the first portion of the first primary data file and the second portion of the second primary data file are synchronized with each other without the rest of the first primary data file and the rest of the second primary data file being respectively synchronized with each other.

15. The storage management system of claim **14**, wherein after the restore operation, the first primary data file and the second primary data file are partially synchronized with each other, based on the backup and restore operations in the storage management system.

16. The storage management system of claim **14**, wherein the storage manager comprises metadata for partially synchronizing the first primary data file and the second primary data file in the storage management system relative to the first portion of the first primary data file and the second portion of the second primary data file, respectively.

17. The storage management system of claim **14**, wherein the first data agent is further configured to:

receive, from the storage manager, first metadata that identifies the first portion of the first primary data file and second metadata that identifies the second portion of the second primary data file; and

wherein the storage manager comprises metadata for partially synchronizing the first primary data file and the second primary data file in the storage management system, based on synchronizing the first portion, but not other portions, of the first primary data file and the second portion, but not other portions, of the second primary data file.

18. The storage management system of claim **14**, wherein the first data agent is associated with a first application that executes on the first computing device and wherein the first application accesses the first primary data file.

19. The storage management system of claim **14**, wherein the first data agent is associated with a first application that executes on the first computing device and wherein the first application accesses the first primary data file; and

wherein the second data agent is associated with a second application that executes on the second computing device and wherein the second application accesses the second primary data file, including the second portion

comprising the restored secondary copy of the first portion of the first primary data file.

20. The storage management system of claim **14**, wherein the second data agent is further configured to:

detect at least one change in the second portion of the 5
second primary data file, and

create, in conjunction with the media agent, a secondary copy of the second portion of the second primary data file; and

wherein the first data agent is further configured to: 10

restore, in conjunction with the media agent, the secondary copy of the second portion to the first portion of the first primary data file on the first data storage device,

thereby transforming the secondary copy of the second portion from a backup format to primary data that is 15

accessible to the first computing device within the first primary data file.

* * * * *